


Aalto University
School of Electrical Engineering
Department of Communications and Networking

Görkem Çakmak

Internet Interconnection Ecosystem in Finland



Thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in Engineering.

Espoo, August 23rd, 2013

Thesis supervisor: Prof. Heikki Hämmäinen

Thesis instructor: MSc. Henna Suomi

Author: Görkem Çakmak**Name of the Thesis:** Internet Interconnection Ecosystem in Finland**Date:** 23.08.2013**Language:** English**Number of pages:** VII + 83**School:** School of Electrical Engineering**Department:** Department of Communications and Networking**Professorship:** Networking Technology**Code:** 38**Supervisor:** Prof. Heikki Hämmäinen**Instructor:** M.Sc. Henna Suomi

For both fixed and mobile network operators, interconnection constitutes an indisputably key element to provide end users with a variety of services. Internet interconnection is particularly an intriguing subject due to the importance of the Internet in our everyday lives and our genuine curiosity to grasp its underlying structure.

This thesis aims to provide a holistic approach to study the Internet interconnections in a nation-centric stance. To accomplish the objective, initially the method that breaks down the key features of the interconnection analysis is introduced. The nation-centric analysis is conducted for Finland by jointly utilizing the Internet registry data and collected Internet routing data. Covering the last decade of the Finnish Internet, the longitudinal analysis yields significant findings for the Internet address usage statistics and the level of multi-homed networks, along with the classification and inference of relationships between stakeholders in the interconnection ecosystem. The implications that the emerging interconnection models pose for the future global service delivery among both fixed and mobile networks are expounded from the perspective of the existing domestic interconnection practices.

The longitudinal interconnectivity study allows us to comprehend both technical and business interfaces between market players by revealing a complete list of customer-provider relationships. Within a national milieu, the assessment of the current Internet market dynamics and future implications of emerging models can be considered in more rationally anticipated manner. Hence, authorities who desire to design new pricing schemes and policies for future networking interconnections can be guided more thoroughly.

Keywords: Internet; Interconnection; BGP (Border Gateway Protocol)

Acknowledgements

This Master's Thesis has been written as a partial fulfillment for the Master of Science degree in the School of Electrical Engineering of Aalto University. The relevant work has been carried out between September 2012 and June 2013 at the Department of Communications and Networking as a part of MoMIE project, funded by the Finnish funding agency TEKES.

I would like to thank my supervisor, Professor Heikki Hämmäinen, for the opportunity to be a part of this project. His profound knowledge and valuable insights were crucial for the embodiment of this thesis. I would also like to thank my instructor, Henna Suomi, for providing me with great ideas and incalculable help throughout the research and writing processes of the thesis during the past year.

I want to thank my friends, as well; they know who they are. Their support was indispensable for me. I would like to thank my dear "*Bal*", for believing in me during all these years of my studies and most importantly for being right by me whenever I needed. Finally, I want to thank my family for their unconditional support and firm beliefs during all my endeavors.

Espoo, August 23rd, 2013

Görkem Çakmak

Table of Contents

ACKNOWLEDGEMENTS	I
TABLE OF CONTENTS.....	II
LIST OF FIGURES.....	IV
LIST OF TABLES.....	V
ABBREVIATIONS.....	VI
1 INTRODUCTION	1
1.1 RESEARCH PROBLEMS AND OBJECTIVES	2
1.2 SCOPE	4
1.3 STRUCTURE.....	5
2 INTERNET INTERCONNECTION ECOSYSTEM	6
2.1 BGP: THE GLUE OF THE INTERNET	6
2.1.1 <i>History of Inter-domain Routing</i>	7
2.1.2 <i>IP Addressing and AS Numbering</i>	8
2.2 INTERCONNECTION ARRANGEMENTS.....	9
2.2.1 <i>Transit</i>	9
2.2.2 <i>Peering</i>	10
2.2.3 <i>Motives to Peer</i>	11
2.3 COST AND VALUE OF INTERCONNECTION	12
2.3.1 <i>Set-up Capital Costs</i>	12
2.3.2 <i>Transaction Capital Costs</i>	13
2.3.3 <i>Operational Costs</i>	14
2.3.4 <i>Peering vs. Transit</i>	16
2.4 INTERCONNECTION MARKET TRENDS	17
2.4.1 <i>Decrement in Transit Prices</i>	17
2.4.2 <i>The Rise of Content</i>	18
2.5 REGULATORY ASPECTS OF INTERNET INTERCONNECTION.....	19
3 MOBILE INTERCONNECTION ECOSYSTEM	21
3.1 CURRENT VOICE AND MOBILE DATA INTERCONNECTION MARKET	21
3.1.1 <i>Voice over IP</i>	22
3.1.2 <i>Charging Models</i>	23
3.2 INTERNATIONAL APPROACHES FOR IP INTERCONNECTION	25
3.2.1 <i>GRX Model</i>	25
3.2.2 <i>IPX Model</i>	27
3.3 REGULATORY ASPECTS OF MNO INTERCONNECTION	30
4 METHODOLOGY.....	31
4.1 WHY NATION-CENTRIC FOCUS?	32
4.2 THE DATA SETS USED FOR THE ANALYSIS.....	34
4.2.1 <i>BGP Data Collectors</i>	34
4.2.2 <i>Content of the Collector Data</i>	35
4.2.3 <i>Supplementary Databases</i>	37
4.3 INTERNET TOPOLOGY STUDIES	38
4.3.1 <i>Inference Method Studies</i>	39
4.3.2 <i>Nation-Centric Approaches</i>	41
5 IP ADDRESS STATISTICS OF THE FINNISH INTERNET	43

5.1	DELEGATION DATA ACCURACY	44
5.2	IP ADDRESS ANALYSIS.....	45
6	A NATION-CENTRIC APPROACH FOR STUDYING INTERNET INTERCONNECTIONS	50
6.1	DERIVING A NATION-CENTRIC INTERNET.....	50
6.1.1	<i>AS-to-Organization Mapping</i>	50
6.1.2	<i>Extracting the Data</i>	51
6.1.3	<i>Stakeholder Categorization and Clustering</i>	52
6.2	AS LINKS: BUSINESS RELATIONSHIPS.....	53
6.2.1	<i>The Inference Approach</i>	54
6.2.2	<i>The Properties of AS Path</i>	55
6.2.3	<i>Observed AS Relationships</i>	57
6.3	RESULTS OF THE LONGITUDINAL ANALYSIS	59
6.3.1	<i>Inference of Observed Links</i>	61
6.3.2	<i>Multihoming</i>	62
6.3.3	<i>Market Trends</i>	63
6.3.4	<i>Public Peering in Finland</i>	64
6.3.5	<i>Peering matrix in FICIX</i>	66
7	MNO INTERCONNECTIONS IN FINLAND	69
8	CONCLUSIONS.....	72
8.1	ASSESSMENT OF RESULTS	72
8.2	EXPLOITATION OF RESULTS	73
8.3	FUTURE RESEARCH	74
	REFERENCES.....	75
	APPENDIX A – THE FLOW CHART FOR EXPANDING THE FILTER LIST OF ASES AND PREFIXES IN LONGITUDINAL ANALYSIS	81
	APPENDIX B – THE HIERARCHICAL LAYOUT TRANSIT LINKS	83

List of Figures

FIGURE 1 - COST OF IXP (PUBLIC PEERING) MODEL [3]	13
FIGURE 2 - PEERING BREAK-EVEN POINT [3]	16
FIGURE 3 - INTERNET TRAFFIC GROWTH VERSUS IP TRANSIT PRICE	17
FIGURE 4 - GRX MODEL [30].....	25
FIGURE 5 - FINANCIAL FLOW IN GRX MODEL	26
FIGURE 6 - IPX MODEL [30]	27
FIGURE 7 - SERVICE TRANSIT MODEL [32]	28
FIGURE 8 - METHODOLOGICAL ACTIVITIES AND STEPS	32
FIGURE 9 - BGP-DERIVED UNIQUE AS PATHS	35
FIGURE 10 - IP ADDRESS ALLOCATIONS FOR FINLAND	46
FIGURE 11 - MATCHED AND FRAGMENTED IP PREFIXES	47
FIGURE 12 - COVERED IP ADDRESSES	47
FIGURE 13 - PREFIX DISTRIBUTION	48
FIGURE 14 - IPV6 PREFIXES IN BGP ANNOUNCEMENTS.....	49
FIGURE 15 - REPRESENTATION OF POSSIBLE BUSINESS RELATIONSHIPS IN AN AS PATH COMPRISING OF FOUR ASEs.....	56
FIGURE 16 - NUMBER OF ACTIVE ASEs.....	59
FIGURE 17 - STAKEHOLDERS OF BGP INTERCONNECTIVITY CATEGORIZED BY SECTORS.....	60
FIGURE 18 - CIRCULAR LAYOUT OF AS CLUSTERS	61
FIGURE 19 - OBSERVED AS LINKS	62
FIGURE 20 - AVERAGE NUMBER OF UPSTREAM PROVIDERS	63
FIGURE 21 - TOP FIVE TRANSIT PROVIDERS OF THE FINNISH INTERNET TRANSIT MARKET.....	64
FIGURE 22 - PUBLICLY EXCHANGE IP TRAFFIC IN FINLAND	65
FIGURE 23 - FICIX PEERING MATRIX	68
FIGURE 24 - DIX INFRASTRUCTURE	70
FIGURE 25 - FLOW CHART FOR LONGITUDINAL ANALYSIS (FIRST PART).....	81
FIGURE 26 - FLOW CHART FOR LONGITUDINAL ANALYSIS (SECOND PART)	82
FIGURE 27 - HIERARCHICAL LAYOUT OF TRANSIT RELATIONSHIPS.....	83

List of Tables

TABLE 1 - SAMPLE OF BGP ANNOUNCEMENT IN MRT FORMAT 36

TABLE 2 - SAMPLE FROM INFERENCE SCREEN.....58



Abbreviations

AMS-IX	Amsterdam Internet Exchange
ARPU	Average Revenue per User
AS	Autonomous System
ASN	Autonomous System Number
BAK	Bill and Keep
BGP	Border Gateway Protocol
CDN	Content Delivery Network
CPNP	Calling Party's Network Pays
DE-CIX	Deutscher Commercial Internet Exchange
DIX	Domestic Internet Exchange
DNS	Domain Name Service
EGP	Exterior Gateway Protocol
ENUM	E.164 Number Mapping
FICIX	Finnish Communication and Internet Exchange
FICORA	Finnish Communications Regulatory Authority
GGP	Gateway-to-Gateway Protocol
GICS	Global Industry Classification Standards
GPRS	General Packet Radio Service
GRX	GPRS Roaming eXchange
GSM	Global System for Mobile Communications
GSMA	GSM Association
IANA	Internet Assigned Numbers Authority
ICT	Information and Communications Technology
IMS	Internet Multimedia Subsystems
IP	Internet Protocol
IPX	IP eXchange
IRR	Internet Routing Registry
ISP	Internet Service Provider
ITU	International Telecommunication Union
IXP	Internet Exchange Point
LINX	London Internet Exchange
LIR	Local Internet Registry
LTE	Long Term Evolution
MED	Multi-Exit Discriminator
MNO	Mobile Network Operator
MPLS	Multiprotocol Label Switching
MRT	Multi-Threaded Routing Toolkit
NAT	Network Address Translation
NGN	Next Generation Network
OTT	Over-The-Top
QoS	Quality of Service
RCS	Rich Communications Services/Suite
RIB	Routing Information Base
RIP	Information Protocol
RIPE NCC	Réseaux IP Européens Network Coordination Centre
RIR	Regional Internet Registries
RIS	Routing Information Service

RRC	Remote Route Collector
SLA	Service Level Agreement
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
ToR	Type of Relationship
TREX	Tampere Region Exchange
UMTS	Universal Mobile Telecommunications System
VLAN	Virtual Local Area Network
VPN	Virtual Private Network



1 Introduction

The importance of interconnection for telecommunications is incontrovertible. One of the basic, yet utmost crucial principles of telecommunications networks is based on the laws that state the value of a network is proportional to the number of connected users and connected networks, such as Metcalfe's law [1]. Whether these laws are overestimations in their quantitative assessments of the value in interconnection, as argued in [2], interconnection is considered to be the most essential element in the development of a competitive marketplace for communications services.

Repeating patterns and historical analogies in the communication technologies – from the telegraph and telephone, to the Internet and mobile devices – have showed that the activity of connecting networks, and thus proliferating via interconnections have been the founding ethos of any point-to-point communications. In the economic context, interconnection is the hidden asset that enables investments channeled into networks yield financial significance, in the long run. In the social context, interconnection is the means that bring each individual user of a communications service into prominence by his/her effect on the overall value for other users as well – the bigger the network, the more valuable it is to both existing and potential members, the phenomenon known as network externalities [1]. In the technical context, interconnection is simply the instrument needed for communications networks to span the globe with technical scalability¹ while not solely exerting the in-house resources and capabilities of a single entity.

Notwithstanding, end users enjoy the seamlessly global and ubiquitous communications mediums, naturally with no or little interest in the happenings on the backstage. Behind the scenes of these mediums (e.g., the Internet, mobile and fixed telephony networks), many separate networks operated by various entities underlie. Coupling the individual networks through contractual agreements, interconnection is simply the glue that holds those systems together.

¹ Scalability refers to the ability of a network to handle a growing amount of work in a capable manner or its ability to be enlarged to accommodate the growth [93].

The interconnection in the communications systems has been interpreted and construed diversely within various terminologies. The interconnection is a compact concept and very well defined within the Internet – the poster child of the Internet Protocol (IP) based networks that thrived in a milieu, exploiting the benefits of open architectures and mildly or none regulated wholesale arrangements. In telecommunications networks, services are subject to strong ex post and ex ante regulations and delivered in a closed fashion with entirely controlled value chains by telecom operators. Thus, technical concepts such as roaming, interworking, signaling, and network management functions surpass the sovereignty of simple interconnectivity notions.

However, regardless of the model and surroundings, interconnections in different communication systems comprise issues that can be classified into three main aspects:

- *Technical aspects* include all the hardware and software arrangements that are necessary to ensure an uninterrupted and seamless flow of communications from one network to another – even if the interconnecting networks comprise of diverse systems and components.
- *Economic aspects* include the commercial bases and policies upon which the parties agree to interconnect their networks and to exchange traffic.
- *Regulatory aspects* consider issues related to the set of regulations that governs the conditions regarding the technical and economic elements of interconnections.

Throughout the thesis, the key components of interconnections are broken down under technical, economic and regulatory aspects, wherein the interconnection related issues in different communication systems can be analyzed most intuitively.

1.1 Research Problems and Objectives

This thesis attempts to answer several interconnection related questions relevant to the Finnish communications market, mainly focusing on the following two:

- Question 1: How to study the Internet interconnections in a nation-centric approach derived from Border Gateway Protocol (BGP) data while considering the issues that are pertaining to both technical and economic aspects of interconnections?
- Question 2: What are the implications that emerging IP interconnections pose for the future global service delivery among both fixed and mobile service providers?

Question 1 points out the lack of nation-centric stance on viable and down-to-earth BGP studies. Inter-domain studies usually have been conducted from either solely technical or economical perspectives. The reason why the interest in studying multidisciplinary aspects of interconnections has been neglected is the disconnection between theory and practice in network economics compared to other industry areas. Operators do use neither theory nor rigorous models; they mostly rely on rules of thumb. On the other hand, academics insist on making models by the help of collected data sets and measurements. When the interconnection related data being sought is neither completely or publicly available, assumptions often replace factual data and become the base for the models that solely interpret a circumscribed and presumptive part of reality. Although the technical aspect of interconnectivity is very well perceived by academics, the backstage of the interconnectivity business, i.e. “the art of peering and transit”, yet remains quite enigmatic for the academia.

Question 2 is a manifestation of the aspiration to analyze the Internet interconnections and Mobile Network Operator (MNO) interconnections together while the evolution towards the ultimate goal of IP based Next Generation Networks (NGNs) propounds emerging approaches. Although these two worlds have markedly different methods and trajectories of development, emerging IP interconnections conceive a possibility for convergence – or quite the contrary, a consummated divergence – between fixed and mobile networks. The exploitation of IP to transport various types of packet traffic in the core and access networks is the most noteworthy driver of change in the current telecommunications industry.

Related to our nation-centric approach, Norton [3] coins a more or less equivalent term, *Regional Internet Peering Ecosystem*: “a community of loosely affiliated network operators that interact and interconnect their networks in various business

relationships within geographical boundaries (country or continent borders)”. The pillar of the prior and most of the present research on the Internet interconnectivity structure is mainly established upon the theme of conventional and superficially global perspectives, missing out locally unique proprieties of Internet regions. In this thesis, an approach to analyzing interconnections is introduced in a nation-centric stance. The main argument here is that each Internet region deserves to be scrutinized elaborately regarding its own characteristics since each interconnectivity ecosystem has its own peering community with their own properties, preferences and obstacles due to the regulations and other geographical constraints.

The initial objective of the thesis is to provide an up-to-date snapshot of all the autonomously administered systems that coalesce to form the Finnish portion of the Internet by the means of BGP. After depicting the current composition of the ecosystem, by adhering to the same methods, the last decade of the Finnish inter-domain construct is investigated with a retroactively longitudinal approach. Additional to the data-based analysis, interviews with interconnection experts, from Information and Communications Technology (ICT) companies in Finland, are conducted to shed light onto the issues of the interconnection market which cannot be studied by merely technical means – e.g., pricing and cost drivers, and the nature of agreements. The view of the national Internet ecosystem is complemented with the identification of MNO interconnections, constituting an all-encompassing view of an interconnection ecosystem.

1.2 Scope

Regarding the scoping of the thesis, we shall draw clear distinctions in several of areas which are liable to become confusing as follows:

- The main focus of the thesis is on IP interconnections accomplished via BGP either for the public Internet or secured and private hub interconnections between MNOs.
- When the pricing related issues are discussed, retail level is left out of scope for this thesis. Although, the linkage between retail and wholesale level prices cannot be ignored completely.

- Access, whereby one network operator uses the facilities of another operator, (e.g. Mobile Virtual Operator Networks) is also out of the scope.
- Direct interconnectivity among service providers by leased line, Virtual Private Network (VPN) connectivity or Internet using IPSec is also out of the scope.

1.3 Structure

The rest of the thesis is structured as follows. Chapter 2 introduces the necessary technical, economic and regulatory background of the Internet interconnections. Chapter 3 summarizes IP interconnection models for MNOs – with emphasis on emerging models. Chapter 4 provides the methodological paradigm and the prior approaches on the interconnection studies. Chapter 5 lays out the results of IP statistics of the Finnish Internet that are natural outcomes of the inter-domain BGP analysis. Chapter 6 introduces the key content of the nation-centric interconnection study approach and presents the results of the longitudinal analysis. Chapter 7 covers the IP interconnections between MNOs in Finland by investigating the deployment process and the current status of the domestic interconnection practices. Finally, Chapter 8 gives the conclusion of the thesis.

2 Internet Interconnection Ecosystem

The interconnection model in the Internet is distinct from the interconnection in telecommunications networks, for reasons that are essential to the Internet architecture itself. Unlike the operator-centric model of the telecom world wherein a strong vertical integration exists, the administration of the constituents and the services offered in the Internet are distributed among varying commercial, non-commercial, and governmental organizations. Due to this fact, the pattern of Internet interconnection agreements and the variety of the participants are discrete from their counterparts in the telecom world.

The fabric of the Internet comprises intricately and dynamically interconnected networks which are achieved through a system of entities called Autonomous Systems (ASes). As the end of 2012, the Internet consists of nearly 42,000 [4] of these decentralized and heterogeneous networks of different business types – e.g., Internet Service Providers (ISPs), ICT companies, large enterprises, educational institutions, governmental organizations, and increasingly content providers. Interconnections between ASes are not only crucial for reachability and accessibility perspectives but also from the performance and quality aspects [5]. The main motive for participating in the BGP interconnectivity is to enhance flexibility and robustness of their networks by the means of interconnection arrangements. After all, how ASes establish and maintain their interconnections determine the way packets are routed, and thus correspondingly impact the quality, choice and extent of the services offered on top of the physical links.

2.1 BGP: The Glue of the Internet

Each and every aspect of this inter-domain routing system, which is a collection of the networks and their administrative organizations' policies, is influenced by the de facto standard inter-domain routing protocol, Border Gateway Protocol (BGP). By default, BGP chooses the shortest route based on policy-compliant next hop, as well as the preference level specified by each AS. These preference levels will be explained in Chapter 4.

Since the commercialization of the Internet – i.e., the transition from the one shared backbone (NSFNET) to the multiple and commercially operated networks – BGP has had a key role of allowing networks to establish interconnectivity without revealing internal and strategic information about their networks. However, in the open nature of the Internet, BGP has certain inadequacies in providing performance and security guarantees which have caused serious instabilities and outages. Several incidents have taken place in recent years due to either malicious or inadvertent misuse of BGP routing [6].

2.1.1 History of Inter-domain Routing

In this section, the evolutionary process of inter-domain protocols is briefly summarized to better comprehend how BGP has turned into the protocol as Internet society harnesses today. More in-depth information about BGP and the history of the inter-domain protocols can be found in [7].

During the days of the ARPANET, the routing protocol evolved into the Gateway-to-Gateway Protocol (GGP), a distance-vector protocol like RIP (Routing Information Protocol). However, it uses a reliable transport mechanism, and unlike RIP, routing updates are sent only when there is a change in reachability status.

In 1984, the Exterior Gateway Protocol (EGP) was introduced. As a routing protocol, EGP is not very advanced. For instance, it does not support topologies with loops in them and needs the network to have a tree structure, in which information flows either up or down (in the direction of either the backbone or stub networks). The main purpose for the protocol was to connect stub gateways – routers connecting to a non-transit network – to the rest of the network and let those stub gateways announce their reachability information.

In 1989, BGP was introduced as a simple path-vector protocol; however, mechanisms were later added to allow each AS to implement their locally defined routing policy, and to keep the policy details to themselves. With the new BGP, routers were no longer let find neighbors on their own; it required manually configuration and running over TCP (Transmission Control Protocol). BGP Version 1 still had the tree structure notion (up, down, or horizontal relationships) as in EGP.

With major changes to the message format, this limitation was abandoned in BGP Version 2. Among many other improvements, a connection collision avoidance method defining how to decide which connection is terminated when two BGP neighbors each initiate a TCP session at the same time was introduced with BGP Version 3. In 1995, BGP Version 4 (RFC 1771 [8]) added Classless Inter-Domain Routing (CIDR), aggregation support, the local preference attribute, and a per-connection hold time. Ever since, BGP has experienced minimal changes [7], including improved route filtering, multiprotocol BGP and application of BGP to other services, such as VPN and Multiprotocol Label Switching (MPLS).

2.1.2 IP Addressing and AS Numbering

From either a technical or business perspective, each public AS is represented by a globally unique AS number (ASN) allocated by the Internet Assigned Numbers Authority (IANA)² or Regional Internet Registries (RIRs)³. Until 2007, ASNs were defined as 16-bit integers, capable of allowing for a maximum of 65536 assignments. Around 73% of these numbers have already been assigned by the end of 2012. With the present assignment rate, the supply of 16-bit ASNs is estimated to be exhausted by 2015. As of 2007, the ASN Registry has expanded to a 32 bit number space which increases the supply of ASNs up to four billion [4].

Each AS advertises at least one block of IP addresses which helps us define an AS in terms of destination. A *prefix* is a portion of IP addresses on the Internet that have the first bits of the addresses in common to signify a set of IP addresses. The number of these common bits depends on the size of the network. Internet routers maintain routing tables that allow them to send traffic to all known public IP addresses defined within network prefixes. BGP is the main enabler of distributing reachability information about these prefixes among ASes.

² IANA is responsible for the global coordination of the DNS (Domain Name Service) Root, IP addressing, and other Internet protocol resources.

³ RIR is an organization that manages the allocation and registration of Internet number resources within a particular region of the world.

However, there are also addresses that are private, and not listed in the BGP routing tables. They might be behind Network Address Translation (NAT) boxes. These private IP addresses are not considered further in this study.

2.2 Interconnection Arrangements

BGP routing decisions are broadly based on routing policies by which routing and interconnectivity intersect with organizations and their strategies. In the early years of the Internet, the nature of interconnections was relatively simplistic, mostly involving ASes with a balanced combination of inbound and outbound traffic. The morphing market conditions and accordingly shifting market power among industrial organizations of the Internet have induced the associated contracts and established interconnections to become more convoluted than commonly understood before [5].

Although the technical issues of interconnections between ASes are very well perceived by academics, unencumbered from a pure technical standpoint, the back stage of the interconnectivity business remains obscure. In order to strive towards obtaining an insightful and thorough perception of the interconnection ecosystem, there are a number of definitions and issues that ought to be comprehended. However fine-grained the interconnection relationships between parties can be, traditionally they present a bifurcated model of arrangements: transit and peering.

2.2.1 Transit

The *Internet transit* is a business relationship whereby an AS sells access to the global Internet [3]. In a traditional customer-supplier arrangement, the customer AS pays the transit provider for transmitting traffic from and to the rest of the Internet. A *transit provider* is usually an ISP the business of which is to provide packet forwarding service for its transit customers. This service is also called upstream transit, Internet connectivity or Internet access. Unlike transit providers, *stub networks* do not provide packet forwarding for other networks. In the hierarchical structure of the global Internet routing, stub networks are at the edges and need transit providers to reach the rest of the Internet. Transit provider ASes may also have their own providers, and are usually represented within different tiers. At the top of the hierarchy, there are about a

dozen *Tier-1* ASes, which do not buy transit from any other AS and connect to each other in a full mesh topology to form the core of the global routing infrastructure [9].

Broadly, multihoming is the main reason for a stub network to involve in AS-level interconnectivity by obtaining a public ASN. *Multihoming* is basically the practice of establishing interconnection to more than one transit provider. A private ASN may be used as well if an AS is only required to communicate via BGP with only one transit provider. Since the routing policy between the private AS and the provider will not be visible in the Internet, a private ASN can be employed for this purpose. IANA has reserved AS64512 through to AS65535 to be used as private ASNs [10].

Typically, the Internet transit service is priced on a per-megabit-per-second (Mbps) basis and metered by using *the 95th percentile* traffic sampling technique. This method formulates a single measurement – the 95th percentile of 5-minute samples – to estimate the transit service volume for calculating monthly fees. In order to boost the use of the transit services, and to be attuned to the competitive environment in the Internet transit market, most providers offer pricing discounts for pre-committing to certain volumes of traffic [3].

2.2.2 Peering

Internet *peering* is a business relationship whereby two ASes reciprocally provide access to each other's customers [3]. When two peering ASes exchange traffic without paying each other, as commonly described, it is a *settlement-free peering*. Unlike Internet transit, a peering arrangement between two ASes does not give either of them access to the Internet via the other. Only, the traffic that is “originated from and destined to the two peering ASes or their downstream customers” is exchanged on a peering link [11]. The traffic from their transit providers or other peering ASes is not allowed. Simply, if a peering arrangement is not settlement-free, then the peering relationship between two ASes is designated as *paid peering*.

The interconnection via peering agreements can be categorized as either public peering or private peering. *Public peering* is usually established at Internet eXchange Points (IXPs), “third-party maintained physical infrastructures” that enable interconnectivity between their member networks [11]. Most IXPs facilitate

interconnection between their members through a shared Layer-2 switching fabric or a Layer-2 cloud. The advantage of these peering fabrics is that with only one router port, a network can interconnect with many other networks while (to some extent) sharing the costs of the links across interconnections with multiple networks. However, it might be troublesome to detect who might be responsible when there occurs a poor end-to-end service with the traffic passing through a multilateral peering point [5].

Private peering is a peering relationship established across a dedicated circuit between exactly two parties, typically via a fiber cross-connection or a Virtual Local Area Network (VLAN) at an IXP or at a co-location center [3]. Private peering provides a higher level of control over bilateral interconnections which facilitates more secure and reliable networks with easier congestion control. Through dedicated resources, private peering helps to identify which network is at fault when there is a problem with the reciprocally exchanged traffic.

2.2.3 Motives to Peer

Internet peering is certainly a delicate but also crucial issue for interconnectivity. Once the requirements, such as balance of traffic flow, geographic reach and market considerations, protected under non-disclosure agreements (NDAs) between two ASes are met, the following rationales are considered to be the motives for establishing peering relationships.

The main incentive of the Internet participants for peering is to reduce operational costs spent on transit. Peering provides a more direct traffic path between the parties while altering and reducing the load on the transit services. When the cost of exchanging traffic with a peering relationship is less than the cost of sending the same amount of traffic through a transit provider, then peering becomes a financially rational option. This will be discussed later on in Section 2.3.4.

By interconnecting directly with peers, an AS can allow its customers to experience lower latency to peered entity's customers, since transit services usually provide a more circuitous path than peering [3]. The liberalization and the consequent growth of the Internet have led to the development of regional networks with substantial

volumes of traffic. Since the commercialization of the Internet in the 1990s, more IXPs have been established in Europe. With IXPs in more countries, there is less need for tromboning⁴ traffic to the US, and a greater portion of traffic is exchanged regionally. Therefore, peering may also allow ASes to have greater control over the routing path and performance of their traffic.

2.3 Cost and Value of Interconnection

There are several cost drivers that are associated with set-up and maintenance of an interconnection. These costs vary widely, depending on the type of the interconnection agreements covered in previous section. Although the costs for interconnecting differ, they can be roughly broken down into two main components: ex ante capital and ex post operational costs.

In this section, several sources – e.g., [3], [5], and [12] – that have attained remarkable attention on the cost and value aspects of the Internet interconnections are exploited to introduce the economics of interconnection in a condensed manner.

2.3.1 Set-up Capital Costs

For peering, the capital costs are well-defined and easy to determine prior to establishing interconnection – even though the costs vary by geographical regions and traffic volumes. The set-up costs are usually fixed, and proportional to the size of the networks. As illustrated in Figure 1, these costs can be categorized as follows:

- Monthly backhaul cost for transmission to the peering point (a fixed-capacity circuit that does not vary with the amount of traffic),
- Co-location costs to maintain the equipment of the networks (e.g., rack space and power),
- Amortized equipment costs (routers, switches, etc.),

⁴ “Tromboning occurs when traffic from one country or region flows through another country to be exchanged and delivered back to the original country or region. In the early years of the Internet, most tromboning took place via the US.” For further information about tromboning see *Analysys Mason’s* report [92].

- Monthly switch port and membership fees at an IXP.

For transit, the only set-up cost to be considered is the interconnection link itself and a router port on each network. The capital set-up costs for transit links are eminently distinctive from peering due to the fact that a transit customer does not require any additional infrastructure outside of its own network. Purchasing transit usually only requires a long-haul capacity to reach the transit provider network. Thus, unlike peering, no co-location or additional external infrastructure is needed. An exception to the simple transit arrangement occurs in the case when an AS extends its network to a distant location where a wider variety of transit providers is available.




FIGURE 1 - COST OF IXP (PUBLIC PEERING) MODEL [3]

2.3.2 Transaction Capital Costs

The transaction costs are more difficult to estimate and vary vastly. As [5] itemizes them in four main categories for a peering agreement:

- Time and effort to contact, and negotiate with a potential peer,
- Configuration of network to support the potential peer's peering policy,
- Engineering resource to support additional network complexity involved in peering vs. simple transit relationships (each time a network adds a link, it increases complexity, and therefore cost of operating the network).

A prior research process takes a significant place in peering due to the fact that there is not only one engineer who understands all aspects of peering or a “peering database” spreadsheet that includes the peering policies of all ASes. A *peering policy* is an articulation of a peering inclination of an AS, expressed either publicly or protected under an NDA. There are essentially four categories of peering policies.

- Open: open to peer with anyone,
- Selective: open to peer with some prerequisites,
- Restrictive: generally not open to new peering,
- No-peering: no intention to ever peer.

The transaction costs are also highly variable on these policies. For example, contacting a network with an open peering policy might be easy, nearly removing all the research cost. However, contacting and negotiating with a larger potential peer which has a restrictive peering policy might be very difficult. Beyond all possible adversity of transaction process, sometimes it might be troublesome to even find the right person in an organization to start the negotiations. Especially, large and medium-sized ISPs usually have selective peering policy when choosing their peers due to the reason that those ISPs may tend to avoid establishing peering relationships with their potential transit customers.

However, for transit agreements the transaction costs are typically much lower since the business relationship is regarded as a typical customer provider relationship where the network selling transit service is inclined to attracting customers and “the customer service provided is one of the benefits that justifies paying more to terminate traffic” [5].

2.3.3 Operational Costs

The essential issue for peering is that the operational costs are supposed to be shared and symmetric, but peers rarely dedicate the same resources to their peering relationship. For instance, when a peering link is degraded or down, peers might not have incentives to troubleshoot the problem promptly since there is no monetary sanction imposed. A similar problem may occur when a link needs to be upgraded.

Thus, a public peering link may become saturated and bring about performance issues if peers do not have the motivation, budget or resources to dedicate.

Aside from troubleshooting duties requiring use of significant resources between peers, there is a never-validated assumption that traffic ratios are correlated to the operational cost of delivering the traffic. When coupled with the de facto industry standard hot potato routing policy, this assumption has a cost-related basis. *Hot-potato* routing means that an AS by default chooses the shortest internal path to a next hop AS network – i.e., it minimizes the number of hops a packet travels in its own network – and thus increases the cost for the receiving network. Assumed that hot-potato routing is applied by both sides, the cost for each will be proportional to the traffic received. Therefore, the balance of traffic ratio between peers implies the balance of cost sharings as well [13].

However, the above-mentioned method does not cohere with the case of Content Delivery Network (CDN) interconnecting with access network ISPs – CDNs will be discussed further in Section 2.4.2 . CDNs can source their content from multiple locations, and normally choose a source close to the destination to reduce latency. CDNs tend to minimize the distance the traffic travels over the receiving access network, which is the opposite of what happens with hot potato routing. Furthermore, imbalanced traffic ratios have inevitably led to disputes over the role of settlements. Especially, the disputes between CDNs and eye-ball-heavy access networks made it loud and clear that settlement-free peering could no longer be sustained when one peer doubles or triples its traffic sent to the other peer [14].

In contrast, the transit provider mostly absorbs the operational costs since Internet transit is offered as a paid-for service from provider to customer. This indicates that operational costs for transit service are asymmetrical, working in the favor of the customer AS purchasing transit. For example, transit contracts usually include Service Level Agreements (SLAs) that have monetary penalties for outages. Hence, a transit customer AS does not have to spend as much engineering resources for troubleshooting outages as with peering. Concerning the performance issues, networks may prefer having the “teeth of a customer-based contract over soft peering assurances that both ASes will work diligently to deal with peering-related issues” [5].

2.3.4 Peering vs. Transit

When the cost drivers and motivations for peering are pondered, ASes agree to peer with each other if both expect to be better off than without peering. In this perspective, for the parties involved, an agreement to peering implies a Pareto improvement [15]. ASes that meet the requirements for peering can choose between peering and purchasing transit while keeping in mind that peering is not a perfect substitute for Internet transit. ASes that do not fulfill requirements for peering must purchase transit or pay to peer; in general, transit can be viewed as a default option. The decision whether to peer or to buy transit is deeply associated with network planning and cost optimization.

In Figure 2, the *peering break-even point* is “the point where the unit cost of peering exactly equals the unit cost for transit”, as Norton defines in [3]. At this point, an AS is in a state of equilibrium between peering and simply sending traffic through a transit provider. To formulate this metric, the monthly costs of peering across the price for transit are calculated while allocating those costs across the capacity of the peering infrastructure. With this graph, the business case for peering can be envisaged to define at which level of traffic exchanged by peering it becomes sensible to peer instead of exclusively purchasing transit from an upstream AS.




FIGURE 2 - PEERING BREAK-EVEN POINT [3]

2.4 Interconnection Market Trends

Highlighting the trends and recent changes in market conditions of the Internet interconnection ecosystem, in this section, issues such as the Internet transit price decline and the rise of content providing networks are discussed.

2.4.1 Decrement in Transit Prices

Over the last decade, the Internet transit prices have significantly decreased due to cost decline of components used in interconnections and competition between transit providers. WIK-Consult [16] reports that unit prices for Internet transit sold to large ISPs and large enterprises have declined at a Compound Annual Growth Rate (CAGR) of -27% during the period 2008 to 2011. Backing up the numbers above, the average price decline expressed in CAGR may be presented from different sources in many ways. However, while providing historical data and speculating about the future, it is crucial to expose the importance of location and traffic volumes in the perception of the transit price decrement.

Among the interconnection hub cities, average price decline differs widely depending on the location. Broadly, price declines are higher where traffic growth rates are larger and number of interconnecting networks abounds. This correlation between the traffic growth and price decline is illustrated by the line “Balanced Demand Growth and Price Declines”, according to TeleGeograph data [17], as presented in Figure 3.




FIGURE 3 - INTERNET TRAFFIC GROWTH VERSUS IP TRANSIT PRICE

Internet transit price CAGR is based on change of median monthly price per Mbps for a fully committed GigE (Gigabit Ethernet) port. As much as geographic variation in prices, port capacity and committed data rates affect the price levels significantly. For example, the lowest 10 GigE port prices have already fallen down to 50 cents per Mbps or less in the Western European interconnection hub cities, while for fully-committed 1 GigE ports transit prices still range from 1.00 dollar to 5.00 per Mbps per month [17].

2.4.2 The Rise of Content

During the past decade, the market structure of the Internet has changed significantly, as a new type of service providers, CDNs have reached their prime. CDNs serve as aggregators of content and provide systems for delivery of traffic directly to the terminating networks [12]. A CDN provides resources to enhance the quality of delivery for Internet content, such as establishing more direct routing via peering agreements, to reduce distance and the number of hops, and caching of content closer to the end users.

A paper on examining changes in Internet inter-domain traffic demands, [18] indicates that CDNs contributed a weighted average percentage of approximately 10% of all Internet inter-domain traffic as of July 2009. In the following year 2010, the total CDN traffic has increased to 20-30% of the traffic on Internet backbones and now estimations are around up-to 45% for 2012 [19]. CDNs have changed the topology of the Internet while flattening its hierarchical structure with providing more direct delivery of traffic and thus disintermediating the role of transit providers. The development of CDNs has induced a consensus that settlement-free peering with traffic-balance requirements should be yielding towards new type of paid peering models between CDNs and access ISPs by which the terms of trade for the various parties along the value chain are reconsidered.

On the other hand, Clark [13] questions whether the CDN market was really one of the two sides of the oligopoly in the Internet, since both access providers and real content owners seem to have even more power than the CDNs. It is pointed out that CDNs are bound to lose the market power and get squeezed in the middle, like transit

provider international carriers, due to their structurally disadvantageous positioning in the evolving Internet ecosystem.

In a similar vein, the distinction between participants of the Internet, such as backbone networks, access networks, and content owners, has started to blur nowadays. It is more intuitive to think of CDN functionality as a business in which many providers with different backgrounds self-provide their content delivery services on varying scales. For instance, Netflix is one of the largest providers of online movies in the United States, and is rapidly expanding into Nordic countries, including Finland. Netflix has been a prominent customer of CDNs to deliver the billions of hours of video every month, until June of 2012, when Netflix announced its own CDN [20].

2.5 Regulatory Aspects of Internet Interconnection

The strong regulations that extensively characterize telephony interconnections have not been applied to Internet interconnections, particularly in the wholesale level. The self-organized and self-regulating nature of the Internet does occur not because the Internet is an “anachronistic, untamed and lawless wild west” environment as [21] figuratively articulates. Conversely, it occurs so because interconnections between ISPs have evolved over decades hand-in-hand with the evolution of the Internet architecture which has shown that self-management is the most effective way to generate and preserve the uniquely meritorious attributes of the Internet. With a hands-off regulatory approach on interconnections, the Internet market has evolved and expanded tremendously throughout the last decades. However, imbalanced traffic ratio problem in peering relationships the parties of which are mainly from different business backgrounds inevitably have recently led to disputes over the role of revenue-neutral settlements.

Especially, the disputes between CDNs and eyeball-heavy access ISPs have made it loud and clear that settlement-free peering could no longer be sustained when one peer doubles or triples its traffic sent to the other peer, and thus creates asymmetrical partaking in cost sharing [13]. Emerging models, such as partial transit and paid peering, represent an aid to cater for a greater diversity of needs and mitigate the drawbacks caused by traffic exchange imbalances.

The deliberations of content market continue in its current form as the debate over network neutrality perpetuates apace. Regulatory authorities are concerned that ISPs might involve in discriminatory practices that may limit end users access to content or applications of their choice. Without the improved transparency into the workings of the Internet ecosystem, regulators may need to interfere with interconnections disputes due to the lack of non-interventional mechanisms to disentangle how ISPs treat their two sorts of interconnections: between their customers and other networks, in particular CDNs [14].



3 Mobile Interconnection Ecosystem

In this chapter, firstly, the current voice and data interconnection markets are discussed. Then, the migration from circuit-switched networks to IP networks and voice over IP concept are covered. Consequently, the wholesale business arrangements that are established between operators are explained. Finally, the chapter is concluded with explaining the existing GSMA (GSM Association) guided interconnection hub models.

3.1 Current Voice and Mobile Data Interconnection Market

Regarding the amount of traffic carried internationally, MNOs are still small players in the IP interconnection market, however, their importance is constantly increasing due to two main reasons. Firstly, worldwide number of subscribers in mobile networks is still increasing, driven by the demand in developing and emerging markets with the low fixed-line penetration and by the increased trend of fixed-mobile substitution in developed countries [22]. Secondly, the rapid uptake of smartphones and the massive roll-out of 3G and 4G networks have resulted in a myriad of new mobile services and platforms.

Cisco data [23] points out a stunning growth for the mobile, in comparison with the fixed Internet; however, the fraction of mobile data traffic in IP networks still remains relatively small. The global mobile data traffic was 2 percent of total IP traffic in 2011 and is expected to reach 10 percent in 2016. The global mobile data traffic grew 70 percent and reached 885 petabytes per month at the end of 2012, up from 520 petabytes per month at the end of 2011. The GSMA [24] also indicates that mobile operator data revenues will eventually overtake voice revenues globally by 2018, as we move towards a fully connected world – i.e., “the Internet of things”. The mobile data explosion is being driven by a surge in demand for connected devices (machine-to-machine communications).

Nonetheless, the growth in mobile broadband Internet subscription and increased usage per broadband subscriber do not automatically generate additional revenues. Both MNOs and ISPs are challenged to develop sustainable business models that

allow them to invest in their interconnections economically while providing novel services to end users.

Unlike the Internet ecosystem incorporating varied type and number of participants, the existing MNO interconnection model consists purely of operators running mobile access networks using GSM, UMTS and nascent LTE air interfaces. While competing, MNOs have a common objective to deliver traffic to each other in a profitable and cost efficient ways. Another objective is to maximize their ubiquity through interworking and roaming agreements for their subscribers to appreciate the full value of increasing number of mobile services.

The existing interconnection architecture is a mixture of direct interconnections between two MNOs – typically used domestically or within a particular region – and indirect interconnections that use an intermediate carrier network to reach the rest of the world for roaming and interworking within “hub” architectures. In Section 3.2, indirect interconnections, established for roaming and interworking between mobile operators, are discussed.

3.1.1 Voice over IP

For decades, circuit-switched voice service has been the primary offering of operators, accounting for the most of the telecommunication network usage and revenue. Today, this pattern has started to deteriorate; voice traffic accounts for only a microscopic share of bandwidth exchanged among networks worldwide. This is true, in part, due to the fact that voice is a low-bandwidth application and plays a much smaller role in the mix of services that end-users consume.

The clear trend for the conventional telecommunications networks, which are predominantly Time Division Multiplexing (TDM) based, is to develop towards IP-based networks. In this sense, rethinking of the interconnections is not a challenging task due to technological changes per se. However, anticipating the economic implications and accommodating regulatory reciprocations to alleviate the disruption that those emerging models may cause will remain to be the compelling steps towards a full-scale IP world.

The transformation from the existing TDM model to IP-based interconnections has the following advantages from operators' point of view [22]:

- Using bandwidth-optimized codec will reduce bandwidth needs, and using IP instead of TDM will ensure better utilization of the point of interconnection. The both improvements will significantly reduce capital expenditure and operational expenditure.
- IP interconnections offer higher flexibility in choosing interconnection partners and architecture.
- MNOs have started launching new multimedia based services such as messaging and presence under the concept of Rich Communication Services (RCS). The interoperability between operators will be one of the critical success factors which will require evolving the interconnections from a pure voice interconnect to a combined voice and multimedia interconnections.

Similarly, voice and video calling services have increasingly been provided by applications that run over the Internet, both by Over-The-Top (OTT) third party service providers and by the traditional networks themselves. The OTT service providers use the network interconnections provided by telecom operators as an enabler for their businesses to bypass traditional international carrier networks. These OTT new entrants have managed to capture 25% of international overall voice traffic volume. Consequently, the increased competition has resulted in a reduction of the voice interconnection price of around 7% per year [17].

3.1.2 Charging Models

In traditional telephony, different charging arrangements can be observed in practice: mainly, *Calling Party's Network Pays* (CPNP), and *Bill and Keep* (BAK) on wholesale levels. Predominantly, service providers adhere to CPNP payment arrangements, where the network of the party that originates a call makes a wholesale payment to the network of the party that terminates the call [25]. Economic theory studies, such as [26], indicate that network operators tend to set these fees “at exceptionally high levels”, and thus they have been generally subject to regulations.

However, a few countries (e.g., the U.S. and Canada) apply alternative arrangements for mobile operators and for non-dominant fixed operators [25]. While freely negotiating termination fees, MNOs are typically subject to the requirement that both network parties of the same agreement must be based on the same per-minute fees. These fees are often set to zero, and therefore it is called *Bill and Keep*.

The wholesale termination pricing and its effect on competition and regulation are broadly investigated in the literature of traditional voice telephony but is scarcely addressed in the concept of data services. As elaborated in [27], a paper on mobile data roaming market, interconnections for the mobile Internet data services do not have entirely the same principles as interconnection agreements between MNOs for voice communications. Interconnections in voice telephony refers to enabling end-to-end users telecommunications traffic, a two way access problem where both providers interconnect to terminate calls and thus, pay each other a termination fee. *Data roaming*, on the other hand, refers to the access of the unilateral service where one operator which does not operate in a respective territory pays to use the entire service of another operator that covers the territory [28].

Regardless of the interconnection charging frameworks, current mobile termination fees are disputed to be too high and argued that they have an apparent impact on the depression of mobile operators' ARPU (Average Revenue Per User) [25]. In the light of these arguments, , the roaming fees for voice calls, texts and Internet access will be eradicated, effective July 2014, under the new telecom regulations in Europe [29]. Although the proposals are highly contentious and receiving stiff opposition from mobile operators, it is believed that MNOs will gain in the longer term while customers will be encouraged to use their mobiles more abroad with reasonable prices – particularly to access the Internet.

With the rapid transition from voice-dominated to mobile broadband networks, MNOs find themselves increasingly squeezed between OTT service providers, such as WhatsApp and Skype, and platform giants, such as Apple, Facebook and Google. In order to counter the threat of becoming transmission pipes that do not capture any value of transactions made using their infrastructure, several new service platforms, the most prominent one being Internet Multimedia Subsystems (IMS), are emerging in favor of MNOs. These platforms will be revisited in Section 3.2.2.

3.2 International Approaches for IP Interconnection

The concept of interconnection in telecommunications networks is relatively elaborated vis-à-vis interconnection model in the Internet with the characteristics, such as reachability, and need for ubiquitous and globally seamless access. The dramatic increase in the type and number of service providers, along with the pre-existence of one-by-one specifically designed solutions, have led to a myriad of complex and fragmented interworking models.

In the light of the global connectedness concept, several international consortiums and forums have been intensively working on defining the landscape for MNO interconnections, such as ETSI TISPAN, i3 Forum, 3GPP and GSMA. The industry momentum has been building around GSMA's interconnection models, and there seems to be a growing desire among these consortiums to avoid divergence. The GSMA is an organization representing the interests of the worldwide mobile communications industry. In the following sections, the models that GSMA endorses and promotes are elaborated.

3.2.1 GRX Model

In this section the interconnections that incorporate only MNOs in the GRX (GPRS Roaming eXchange) model, merits of which are under consideration by the GSMA, are expounded. The GRX is an exchange model established to support for the interconnection of roaming and interworking networks associated with the GSMA. [30] The architecture of GRX model is illustrated in Figure 4.




FIGURE 4 - GRX MODEL [30]

Since 2000, with the advent of 2.5G technology, GSM operators have been using the GRX network to route IP-based roaming traffic between visited and home operators. In particular, the GRX is used to support traffic applications including: GPRS and 3G data roaming, LTE data roaming, WLAN roaming, interworking of messaging services and IMS. While there are some SLA associated with GRX, there is limited Quality of Service (QoS). In some cases, MNOs are connected in bilateral arrangements to each other when it makes economic sense. However, the key feature of GRX is its ability to employ a hub architecture which can enable establishing only one connection to reach multiple MNOs. Thus, the model reduces the need to establish dedicated links between each operator to support mobile data roaming.

The CPNP is also the base payment method between MNOs in the GRX model, however, charging between GRX providers and MNOs are volume or capacity based monthly payments, similar to the Internet transit. Exchange of IP traffic with no monetary exchange in BAK type of agreement, similar to the model used in settlement-free peering contracts in the Internet, can be also seen between GRX providers as illustrated in Figure 5.




FIGURE 5 - FINANCIAL FLOW IN GRX MODEL

Based on the industry specifications and recommendations developed by the GSMA, MNOs are interconnected with GRX providers via dedicated IP interconnectivity (leased line circuits), or logically separated connections from the public internet while aiming to establish a “mobile Internet”. GRX providers have roughly the same role

for MNOs as Internet transit providers have for ISPs in the Internet. The technical aspects in GRX are also excessively similar to the Internet – including common routing protocols (BGP), same AS numbering and IP addressing and look-up mechanisms (DNS), and even the very same physical locations for interconnections – e.g., the Amsterdam Internet Exchange (AMS-IX) and Ashburn Equinix [31].

3.2.2 IPX Model

As we move to a full-scale IP interconnectivity, an increasing demand in quality for end-to-end applications that are delay sensitive and conversational in nature – e.g., voice and video calling – has shaped. Consequently, there has formed a prospect for a new interconnection methodology that adds QoS and service aware capabilities, as well as usage-based charging. To cater these needs, GSMA has developed the IP eXchange (IPX) built upon the architecture of the GRX while introducing end-to-end QoS and a number of non-GSM, new stakeholders – such as, Fixed Network Operators, ISPs and Application Service Providers and Content Providers. Hereafter, to address all these participants collectively, *IPX service provider* term is used. The IPX architecture is illustrated in Figure 6.




FIGURE 6 - IPX MODEL [30]

As QoS is a cornerstone of the IPX specifications, an IPX network – as the GSMA specifications define – must only use private IP interconnections, logically isolated from other networks accessible from the Internet even though the addressing is handled by using public IP addresses [30]. IPX represents a billing model as much as

technical model. It enables cascading billing; a chain of billing can be established from application and through each IPX hub to end users. IPX supports three types of interconnectivity options, commercial models of which are discussed in [32] and [33] – the latter mainly focuses on voice over IPX services.

In *Transport Only* bilateral agreement option, IPX service providers use transport layer while cascading of responsibilities such as QoS persists but not cascading of payments. Each service provider pays their respective IPX provider for the transport capacity and optional termination charges according to a bilateral agreement between service providers. The exchange of traffic between two IPX providers can be done on a paid model or a free peering basis as in the free peering model between GRX providers [33].

With *Service Transit* bilateral connection, two IPX service providers use service layer and the transport layer with guaranteed QoS end-to-end. In addition to capacity charges, service providers can optionally pay directly termination costs according to bilateral agreement (settled outside of IPX) or payments can be cascaded within the chain between sending and receiving party through IPX providers. Figure 7 depicts this model with financial flow between interconnection participants.

The last connectivity option is *Service Hub* model that provides multilateral interconnection with guaranteed end-to-end QoS and service-based charging. IP traffic can be routed from one participant to many destinations via a single agreement with an IPX provider. This option is expected to be the preferred long term option [34] .




FIGURE 7 - SERVICE TRANSIT MODEL [32]

The crucial advantage of these models lies under the fact that the IPX ecosystem establishes *a second tier of Internet* that enables the concept of equitable payments between all stakeholders in the value chain while ensuring the quality guaranteed service delivery – which is vital to the success of time-sensitive applications. While the CPNP charging model constitutes the core aspect of the IPX business framework, the GSMA recommendations [32] suggest that IPX providers will be able to support a variety of charging principles built on top of this core or alternative models – such as session based, data volume based, event based– typically varying on a per service basis.

Undoubtedly, the service-based interconnection charges over the volume-based model will prevail for services that convey premium quality and experience to end users. Although CPNP model is one step ahead of its alternatives in terms of efficiency, there is no single “one-size-fits-all” approach in interconnection charging models maximize economic efficiency in all circumstances [35]. With the emergence of new interconnection opportunities and the advent of a wide range of new services, the need for coherent charging mechanisms that can correlate efficient retail pricing of end-user services with wholesale charges assuring equitable distribution of network cost will increase in the future.

In order to facilitate the introduction of commercial IMS services, mobile operators are taking collaborative action to develop Rich Communications Services (RCS) platform. The idea is to leverage the unique operator proposition of universality and QoS within an objective of transitioning traditional voice and messaging services into an all IP and LTE world. By entitling operators to play the role of service providers as well as platform providers, RCS and Voice over Long Term Evolution (VoLTE) are considered to be the driving forces in MNOs’ plan to accelerate the proliferation of service-aware hub interconnections [36].

On the other side of the story, third party service providers can be prevented from MNO interconnections over IPX platforms unless third parties participate in the IPX ecosystem and negotiate MNOs’ pricing and quality arrangements. Nevertheless, this might not be as simplistic as it appears; blocking or allowing access is not a twofold dilemma. The services offered by OTT players are, in most cases, not conspicuously eroding operator’s business but are often contributing on mobile Internet usage per

user. At this stage, MNOs will most likely adopt an OTT mentality and moreover partner up with third party players to harness the advantages of emerging models that could be facilitated solely by neither the Internet nor operator models. It is an impractical endeavor to form a single service platform offers superior flexibility to both service providers and end users [37]. Openness and compatibility between platforms, even with hybrid solutions, can influence the level of received benefit for both end users and service providers, and hence maximize the functionality and efficiency of the emerging IP interconnection models.

3.3 Regulatory Aspects of MNO Interconnection

MNO interconnections have always been subject to strong regulations. Market has evolved in a manner that oligopolistic industry has become a natural outcome of mobile business. The gradual demise of circuit-switched voice traffic and prominent raise of mobile data are creating delicate challenges for regulatory authorities.

The crux of the evolution towards future networking lies in reconciling between MNO and ISP models. IPX, here, has a prominent potential to sit in-between and fill the gap. Countervailing the pitfalls of each model, IPX model is expected to ensure revenues generated by real-time IP sessions continue to flow via the networks while allowing MNOs and ISPs to maintain their position in the value chain. However, the part that may concern regulators is that this emerging model might be the first indicator of the embodiment for the foundational separation of the Internet, in terms of QoS. Reference [38] argues the tendency of ISPs and MNOs to use the QoS mechanisms in a closed manner to guarantee the compensation of upgrade cost for deploying QoS and enhance revenue opportunities while jeopardizing the open nature of the Internet with creating opportunities for vertical integration.

The ability of the Internet model for interconnection agreements, which has been inherently contradictory with telephony model, to produce efficient results and disrupt telecom models is broadly argued in favor of allowing full-scale IP world. The current regulatory inclination in the telecom industry indicates that the abnormalities in the termination fees have to be reconsidered. The inherited pricing models should be attuned to accommodate a healthier and smoother transition to IP interconnection infrastructures and to cope with the competition from OTT service providers and platforms by reflecting retail level needs on the wholesale termination fees.

4 Methodology

The ultimate goal towards an all IP communications environment brings out a stimulus for an emerging shift in the status quo perception of IP interconnection practices. Despite the frequently-uttered notion of “convergence” in the literature, the Internet and telecommunications worlds have markedly different stances towards the interconnection concept as we covered in the previous chapters. Therefore, the means by which these stances can be studied ought to be different, as well. This chapter sets the scene for how interconnections in the Internet and telecom ecosystems can be studied and further scrutinized.

Unlike the telecom world, many openly available data sources exist for studying Internet interconnectivity. There are three main types of data sets that have been available for researchers: a) BGP tables, b) Internet Routing Registry (IRR) information, and c) traceroute measurements. BGP tables have been collected by the University of Oregon’s Route Views project [39] mostly in the US and Routing Information Service (RIS) project [40] run by Réseaux IP Européens Network Coordination Centre⁵ (RIPE NCC), mostly in Europe. Meanwhile, traceroute-based datasets have been gathered by CAIDA [41], by an EU project called Dimes [42] and more recently by iPlane project from University of Washington [43]. The Internet routing and topology studies have been harnessing monitored routing announcements and trace-route derived data to conduct inter-domain (AS-level) constructs of the Internet – along with finer granularities such as point of presence level and prefix level topologies.

On the other hand, telecom industry is a closed world, and very little information exists in the literature when it comes to the inter-operator interconnections. In the absence of open source data, the market reports from industry researchers and regulators are, indeed, useful. Interviews with industry experts also provide a genuine

⁵ Although similar in name, the RIPE NCC and RIPE (Réseaux IP Européens) are separate entities. The RIPE community refers collectively to any individual or organization that has an interest in the way the Internet is managed, structured or governed. The RIPE NCC provides administrative support to RIPE. [94]

way to understand the real market dynamics and concerns, especially for grasping an insight of MNO interconnections which cannot be investigated by solely relying on and analyzing technical means.

While striving towards seeking answers to the research questions listed in Chapter 1, there are many aspects and subgoals – some of which are orthogonal, while some are interdependent – that should be conceived, in a systematic manner, with a series of well-defined steps in order to yield consistent results. The embodiment process of the method has been an ever-evolving outcome in practice towards the aforementioned objectives. Therefore, some of the results will be a couple of times exploited for explaining the content of the method, as well as boosting the validity of the method itself. In Figure 8, a flow chart of methodological steps and final outcomes is depicted.




FIGURE 8 - METHODOLOGICAL ACTIVITIES AND STEPS

4.1 Why Nation-Centric Focus?

The reason for defining the scope tailored to national level is that the diversity of complex and dynamic global Internet ecosystem simply prevents us from conducting generalized models regardless of the local surroundings. Context – the environment, the participants and the dynamics that are at the core of any recursive fragment of the Internet – is significant. Here, the Internet region abstraction can help us to compartmentalize the whole Internet context into a bounded portion where ASes need

to adjust themselves and operate according to the regional features if they so wish to participate in that particular Internet ecosystem.

Scrutinizing the global Internet on the inter-domain level is often too coarse-grained to observe relevant interdependencies between the participants of a particular Internet region. More or less each Internet region presents a predictably similar set of operations performed by a common categorization of participants: ISPs, stub networks, content providers, IXPs and a regulatory authority in their regional or inter-regional habitats.

Regionalizing the Internet also requires defining borders between these regions. Even though country borders do not practically apply interconnection-wise in the Internet; yet, using the term Internet region as a portion of the Internet contained within the boundaries of a country makes better sense to refer to the analogous set of diversely interconnected market players that together constitute an autonomous Internet ecosystem. For counter-arguments of the Internet region approach, it is worth bearing in mind that the Internet has never been “a whole” to begin with; it has been an assemblage of many components – i.e., a heap of pebbles rather than a monolith.

Likewise, as IP-based backhaul interconnections evolve, different models of interconnection between MNOs are being shaped both at national and international levels. Activities on the international level are developing in a generally more organized fashion as explained in Section 3.2. These well-organized interconnection activities are led by a number of industry organizations among which an industrial momentum has already been built around with a good deal of commonality and a desire to avoid divergent solutions.

By contrast, only a few national regulators are actively defining the future interconnection landscapes in their countries. However, these countries are taking different approaches within different timelines and no harmonization in their varying actions. Besides, some countries are just forbearing from engaging in any activities in this area at the moment. Due to the fact that MNO interconnections are associated with numerous regulatory concerns and nation-specific conditions, a national level focus on MNO interconnections does also make better sense, once the structures of international interconnections are comprehended.

4.2 The Data Sets Used for the Analysis

The methodology and the scrutiny heavily avail themselves of BGP-derived data from RIS route collecting project and of several other databases which will be explained in Section 4.2.3. The latter category of databases has been harnessed as complementary to the main BGP data source.

4.2.1 BGP Data Collectors

With an explicitly-stated purpose declared [44], as a service to the Internet community, in the late '90s two organizations started collecting and providing real-time BGP routing data gathered from a number of international backbone networks. These two projects, namely RouteViews and RIS, host multiple data collectors officially called as *Remote Route Collectors* (RRCs). These collectors establish BGP sessions with operational routers in ASes that the routing information is collected from. Throughout the thesis, we will call each operational router connected to a RRC a *monitor* and the AS that the router belongs to a *monitor AS*.

Both projects are motivated by the interest of operators for determining how the global routing system would view their prefixes and ASes. The data collected by these projects has been an important asset to network operators for debugging purposes and to the academia for Internet topology studies. Typically, these collectors record all the BGP tables and updated announcements that they receive from the neighboring ASes over time. RRCs are mainly connected to large ISPs and located at large IXPs. Both RIS and RouteViews projects have collected data from several hundreds of ASes from different vantage points of the Internet for more than a decade, and the collected data sets have been made publicly available in open data formats.

Having aggregated and parsed crude data of BGP routing tables from RRCs of RIS project, we are able to monitor ASes and prefixes that are relevant for the Finnish Internet from various vantage points. RIS project includes overall 17 RRCs that collect and store Internet routing data from their approximately 400 peers. These collectors provide us their view of the Internet from different perspectives around the globe.

Nonetheless, a collector can only see what the monitors (connected routers) choose to send along. Contrary to common belief, one does not observe the Internet as seen by monitor routers. It is more probable to expect to “anticipate what a downstream neighbor of the monitor router might receive” [45]. In order to distinguish the view that these collectors provide, we will use the term *public view* to define the interconnectivity structure inferred from publicly available BGP data sources. In Figure 9, the distribution of BGP-derived unique AS paths to Finnish prefixes and ASes – collected and parsed from each RRC for each year of the analysis – is presented.




FIGURE 9 - BGP-DERIVED UNIQUE AS PATHS

4.2.2 Content of the Collector Data

The data in RIS project comprises of Routing Information Base (RIB) entries – collected in every 8 hours separately from each collector — that are encoded in a distinct sequential Multi-Threaded Routing Toolkit (MRT) record. This format is described lengthily in RFC 6396 [46]. In Table 1, a BGP message sample in MRT format is presented. The BGP messages in RIBs provide us with several significant

attributes and information of pertinent ASes. These attributes are notably used to determine the best possible route to a prefix when multiple paths exist to reach this prefix [47]. The detailed information on attributes can be found in RFC 4271 [48]. Among these attributes, AS Path, Local Preferences, Multi-Exit Discriminator (MED) and BGP Communities values between neighboring ASes are particularly investigated throughout the analysis.

BGP Type	Time (UTC)	Peer IP	Prefix	AS_PATH	Local Pref	MED	Community
TABLEDUMP2	1359057613	195.22.216.18	195.140.192.0/22	3356 6667 719	0	100	6667:3000

TABLE 1 - SAMPLE OF BGP ANNOUNCEMENT IN MRT FORMAT

- *AS Path* is an ordered list of ASes that the route advertisement has traversed. Table 1 shows an example in which a route is passing through three ASes. For studying AS-level interconnections, researchers have focused on AS Path attribute with the utmost interest.
- The *Local Preference* (Local Pref) attribute is used to prefer an exit point from the local AS. Local Preference is a setting for the local AS and only passed to iBGP (internal BGP) peers, not passed to eBGP (external BGP) peers. Thus, in BGP announcements collected from RRCs, the Local Pref is always marked null. However, the Local Preference used in a local AS still can be extrapolated by the help of BGP communities attribute.
- The *MED* is basically a suggestion metric to an external AS regarding the preferred route into the AS that is advertising the metric. Commonly used for mitigating (moderate) traffic imbalances between peering ASes, MED allows an AS to notify a neighbor AS of its preference as to which of several links are preferred for inbound traffic.
- The *BGP Communities* attribute is used by AS administrators to define and group destinations which share some common property. This attribute has several usage purposes. For instance, it might be applied in multihoming routing, as defined in RFC 1998 [49]. By employing the AS-based customization of the Local Preference attribute, it be can used to adjust which routes should be preferred primarily over routes advertised for a backup link. Standard values of BGP communities are described in RFC 1997 [50].

4.2.3 Supplementary Databases

Each Regional Internet Registry (RIR) has its own database, a significant part of which is used for aggregating routing information. The regional registries are AfriNIC for Africa, ARIN for the North America, APNIC for Asia and Australia, LACNIC for Latin America, and RIPE NCC for Europe, the Middle East, and Central Asia.

The RIPE Internet Routing Registry (IRR) database contains registries by which network operators are able to publish their routing policies and their routing announcements. The IRR database is used by network operators to discover peering agreements, determine optimal policies, and more recently, to configure their routers and filters accordingly. The data is accessible by a WHOIS query (whois.ripe.net) and through a web interface [51]. RIPE IRR data is utilized for the AS-to-organization mapping analysis, as explained in Section 6.1.1.

The RIRs' databases also contain information about allocations and assignments of IP address space, which is quite often referred to as *RIR delegation data*. The delegation data contains ISO 3166-1 2-digit country codes for each block specifying the country of the allocation [52]. However, no specified rules are defined for this value to indicate the country where the addresses are used. IP address geolocation based on the delegated data cannot be always accurate, since large IP address blocks assigned to one country may belong to large international organizations that are widely spread over multiple countries. In Section 5.1, this issue is revisited with an evaluation of the delegation data accuracy (geolocation-wise).

Another additional information that is made use of is the list of Local Internet Registries (LIRs), a term used to describe the members of RIPE NCC. LIRs are responsible for the distribution of address space and registration of the address space on a national level. LIRs also ensure that policies and procedures are followed on the local level. The LIR organizations are mainly ISPs that assign and allocate address space to their customers. Currently, there are 367 member LIRs offering service in Finland, however, 214 of these registries (58%) are based in other countries than Finland.

There are also a number of sources that maintain mainly IXP related information – e.g., list of participants in IXPs, traffic statistics and subnet prefixes used in layer-2

clouds. Peeringdb [53], Packet Clearing House (PCH) [54] and Euro-IX [55] are the websites used for collecting IXP related data in this study.

4.3 Internet Topology Studies

The analysis of ASes and their relationships has been an attractive research area over a decade, starting with Govindan and Reddy [56]. Another well-known and excessively cited work of Faloutsos brothers [57] propounds the inter-domain view of the Internet can be readily and accurately obtained from the available public view BGP data. This approach has been followed by an influx of significant, but heavily graph theory oriented research activities in this area [45].

Trying to lay down a precise interpretation of the use of “Internet graph”, in the existing literature, the term has been used to refer to a virtual construct created via BGP. It is this particular interconnectivity structure that is focused on, in this thesis. Interconnectivity between ASes may allude that ASes have physical interconnectivity, established BGP sessions to exchange data traffic, or are in a business relationship protected under an NDA. All of these implications for Internet interconnectivity are reasonable, however none are equivalent.

An important aspect of the BGP is the hiding of the internal structures; “BGP allows networks to exchange routing information between them without revealing strategic information about their own networks” [45]. Hence, stakeholders of the AS-level Internet tend to view their BGP interconnectivity as proprietary information. ASes are often averse to reveal sensitive business information, e.g., the number of routers inside their networks, the geolocation and topological structure of their networks, the list of transit customers and peers. Although it is outside the scope of this thesis, we shall dwell on addressing the current deficiencies on measurement and inferences of the AS construct to define what is possible (measurability-wise), for interconnectivity analysis.

The crucial problem has been the imperceptive reliance on public view BGP data as the sole source of information. By its nature, “BGP is an information-hiding rather than information-revealing routing protocol” [45], and utilizing it for mapping the complete Internet topology is not a purposefully executed measurement method, to

begin with. Similar to the BGP data, traceroute measurements and looking glass servers are also debugging tools that were not mainly intended to reveal topologies. Nonetheless, the fact that traceroute provides a router path while BGP returns a path in adjacent AS-hops indicates that these measurements are orthogonal and hence can be used as complementary methods. IRR data provides useful information in general but is known to contain a significant amount of out-of-date and incomplete data due to the voluntary nature of registry entries. Unfortunately, these three methods are more of a reflection of what is measureable than what is supposed to be measured to make proper interpretation of the Internet's complete AS-level interconnectivity structure.

Moreover, the common method of abstracting ASes to generic and identical nodes without any internal and economic structure is realized to be an over-simplification that limits our ability to interpret the rich content in the inter-AS relationships. The traditional graphic presentation of the Internet topology and the prior works' insistence on abstracting AS-graph view of the Internet are the main reasons that have caused the disparity between the trajectories of the academia and the industry-oriented research on this subject.

4.3.1 Inference Method Studies

The AS-level construct of the Internet is not effortlessly available due to the decentralized nature of the Internet. Even though it is rather easy to collect a roughly complete set of active ASes, it has been proven to be a difficult task to collect the complete list of inter-AS links. There has been a great effort in the research field of Internet topology, mainly focusing on business relationships between ASes. As mentioned above, ISPs consider the details of their business relationships as proprietary information and tend not to reveal them. Therefore, researchers have been relying on AS relationship inference algorithms in order to depict a view of Internet business interfaces. During the last decade, researchers have introduced a number of algorithms to infer AS relationships. However, these algorithms have produced conflicting results, since inference algorithms are limited by the fact that they rely on heuristic assumptions when AS interconnectivity information is not sufficient.

Gao's seminal work [58] has inspired many researchers to discover approaches and algorithms to infer AS business relationships by using publicly available BGP data.

Gao's assumption indicates that each AS path must comply with the following hierarchical routing pattern: "an uphill segment of c2p (customer-to-provider) or s2s (sibling-to-sibling) links, followed by zero or one p2p (peer-to-peer) links, followed by a downhill segment of zero or more p2c or s2s links" [59]. AS paths with this hierarchical structure are called valley-free or simply valid paths. The paths that do not follow this structure are called invalid and may derive from BGP misconfigurations or from complex BGP policies that do not distinctly fall into the simple classification (c2p, p2p, s2s). The valid paths concept will be revisited in Section 6.2.

Subramanian et al. [60] provided a formulation based on the concept of valid paths, but simplified the problem by excluding the inference of s2s links – relationship between ASes belonging to the same organization. The authors investigate the inference of business relationships with the Type of Relationship (ToR), a combinatorial optimization problem. The approach determines a rank for each AS and this rank is used to measure the closeness of an AS to the graph core. The heuristic accordingly infers relationships by comparing ranks of adjacent ASes. For example, if the ranks are about similar, the algorithm classifies the link as p2p, otherwise as c2p.

References [61] and [62] independently developed mathematically approximate solutions to the ToR problem. They proved that it is not possible to infer p2p relationships under the ToR formulation, and thus their solutions infer only c2p relationships and ignore p2p and s2s relationships. In [63] and [59], the authors identified several other issues, like improved algorithms that determine not only c2p but also p2p links for those can be detected from BGP data. Reference [64] introduced the idea that the resulting graph should be acyclic – the valid paths should contain no cycles – and presented a new algorithm that does the assignment and reduces the number of cycles. All these improvements have achieved more accurate AS relationship inferences, however still yielding contradicting results with each other.

A utilization of additional information sources in BGP announcements for inference methods and for broader networking purposes by extracting BGP communities and local preferences attributes – rather than the sole usage of AS connectivity information – is proposed in [47]. BGP communities attribute provides a scheme for grouping AS destinations into separate entities in which similar routing decisions may

apply. Although usage of the attribute is optional, it has become intensively used by operators to facilitate flexible routing policies. Communities and their parameters are not standardized; many ASes explain the meaning of their communities values in IRR or on their own websites. A large portion of communities are outbound communities, thus provide crucial information – e.g., type of business relationship, geographic location, local preferences – about their adjacent ASes.

4.3.2 Nation-Centric Approaches

The challenging question for nation-centric approaches is how to define the participants elaborately when considering a country-wise portion of the Internet. A regulation-oriented study, [65] approaches the question with a method that maps nationally relevant networks of ASes by identifying a smaller set of ASes acting as “points of control” for the rest of the national Internet. Their methods, analogous to ours, begins with assigning each AS to a country, however in contrast with ours, continues with solely relying on the IRR data for defining AS that plays central role for a nation-centric analysis of the Internet. Registry data maintains the authoritative list of ASNs and IP blocks associated with each country included in their own regions, albeit, reality does not always align with the data provided from these registries.

Another work on nation-state understanding of the Internet, [66] presents (active) measurement of the Chinese Internet AS graph based on traceroute data probed from servers of major ISPs inside mainland China. Their obtained Chinese AS-level Internet graph is a small regional sub-graph of the global Internet. It is claimed to contain only ASes from the mainland of China. However, the stage of categorizing and filtering Chinese ASes, and the impact of international players on the national Internet ecosystem remain rather destitute of a rational explanation.

The most well-rounded study on nation-centric approach, [67] is also determined to define the set of ASes that compose the nation-centric part of the Internet for Germany. In their framework, they start with extracting IP-blocks, which are either registered in RIR with specific country code or provided with additional information of the administrator entities in IRR database. Additional to the registry information, the approach in [67] aims to foster the validity of the data accuracy by also going

through different data sources like Cymru [68] – a commercial IP address geolocation provider – and the RRC of RIS project located in Deutscher Commercial Internet Exchange (DE-CIX), in Frankfurt, whence the most relevant BGP data for Germany can be collected.



5 IP Address Statistics of the Finnish Internet

The prefix reachability information of the Internet destinations is propagated through BGP route announcements. Originated from the AS that the prefix belongs to, these announcements are selectively propagated to other ASes in accordance with routing policies. A *global routing table* lists every single prefix on the Internet. Each AS will have a different global routing table, since each AS will have different AS-level paths to each prefix [5]. This property helps inferring the AS-level connectivity, as much as providing nonpareil information about the dynamics of prefix usage. By correlating the IP allocation data with actual announced prefixes, how efficiently ISPs and other stakeholders in the Internet use their allocated IP addresses can be examined.

The Internet has been experiencing a tremendous growth over the last decade. The nature of the Internet access has also evolved from being reliant on fixed access toward mobile access through smartphones and tablets. Correlatively, the Finnish Internet market has continued to be characterized by a highly innovative and competitive mobile market with increased fixed-to-mobile substitution. The increasing number of access devices, the surge in mobile and fixed broadband and the proliferation of end-to-end services point out for an increasing demand for more IP addresses in the Internet.

As a natural consequence of this expansion of the Internet, the global routing table has expanded enormously with new allocations, fragmentation and finer segmentation of IP addresses. An evidence of this growth is the size of global BGP routing table; the collected data shows that the routing table size has nearly multiplied by 10 times over the last decade to reach all prefixes in Finland. The main reason for this multiplication in size is the tendency of operators to subdivide allocated IP blocks into several individual prefixes and announce them separately. Another reason is the various traffic engineering techniques, such as traffic balancing and multihoming, which lead announcing the same IP address blocks covered by more than one prefixes.

5.1 Delegation Data Accuracy

For the case of Finland, there are 204 ASes registered in RIPE (the relevant RIR for Finland) database with “FI” country code attribute – country code is a mandatory attribute for each allocated AS number and IP address block. However, the BGP study reveals that 46 (23%) of these ASes are actually not routable, hence not reachable in the Internet, as neither do they appear in global BGP routing tables nor appear to be used by organizations from other countries. Besides including a number of dormant ASes, single-handedly usage of the registry data may lead to spurious country assignments. For example, our analysis points out that there are eleven ASes that belong to two major Finnish operators (Elisa and Telia Sonera) and to several other Finnish networks are registered with “EU” country code. The last but not the least, the operator that has the highest number of downstream customers in Finland is a subsidiary of a company business operations of which have spread in other Nordic countries. Therefore, it does not have a “FI” country code but plays a crucial role in the Finnish Internet.

Apparently, AS-level approach is sometimes too course-grained to identify country-specific actives of multinational operators. It is impractical to toss out the role of these multinational players. Most of the time, these ASes are, in fact, the gateways of the national Internet that provides the local ASes – which do not have numerous opportunities to interconnect with many ASes and enlarge their interconnectivity portfolio – with the global routing tables.

In this part of the study, the accuracy of the IP allocation data for IP address geolocation is investigated, regarding the Finnish portion of the Internet. A related research comparing geolocation accuracy of many sources, [69] shows that country identification from RIR delegation data disagrees with the commercial geolocation databases, in average for 4.4% of all allocated addresses. This actually suggests that, contrary to the common presumptions, “prevailing majority of addresses are typically being used in the country which they were delegated to”. The authors of [69] indicate that commercial geolocation service providers generally agree on IP-address-to-country mappings, however, geolocating on the city level has remained to yield conflicting results.

Thus, the IP address allocation data for Finland is compared with a (presumably) more accurate IP geolocation database of a commercial provider, MaxMind's GeoLite [70]. This database has been claimed to obtain the highest level of agreement (99.1%) among other commonly used commercial databases for country geolocations [69]. As of November 2012, IP allocation data of RIPE indicates that there are 13,613,952 IPv4 addresses allocated to Finland with "FI" country code, whereas, according to MaxMind's GeoLite database, there are 13,755,886 IPv4 addresses geolocated in Finland. When the IP addresses in both databases are cross-checked, a difference in the country mapping for 1.28 % of the IP addresses was observed. 176,879 IP addresses from GeoLite database are claimed to be used in Finland but assigned with different country codes. On the other hand, the difference is not very large since the mapping was identical for 98.72 % of the allocated IPv4 address space. Furthermore, there are IP addresses allocated with "EU" country code that are announced from the aforementioned major Finnish ASes. When those announced prefixes are also taken into account, the discrepancy can be minimized down to 0.33%, meaning that there are only 45,296 addresses that do not coincide with delegation information provided by the relevant RIR.

The reason for being keen on winnowing the delegation data and wielding which data source and to what extent can be treated accurate is that RIR delegation datasets are historically retraceable and freely accessible. The country information provided from registries is an indispensable element as a starting point especially for retrospective analysis, even though the data can be stagnant and has to be validated by supplementary databases. In order to reduce the dependency to commercial datasets, the base information from authoritative organizations should be highly utilized.

5.2 IP Address Analysis

In this part of the thesis, for each year of the last decade, BGP routing tables are parsed to investigate the IPv4 address allocation records of Finland – as previously referred to as delegation data– and their impact on the BGP routing table size. The number of used and unused allocations, and the percentage of prefixes advertised as identical or fragmented are the main focus of the delegation data analysis.

The RIRs are responsible for allocating IP address blocks to ISPs. Consecutively, ISPs assign IP address blocks from their allocated addresses to their customers: enterprises and individual end users. Allocated IP address blocks, represented by prefixes, are utilized if and only when the prefixes are advertised into the global routing system, otherwise they remain dormant. Figure 10 represents IP addresses that are allocated but appear to be not in use, and IP addresses that are being used and reachable actively from collected BGP routing tables.




FIGURE 10 - IP ADDRESS ALLOCATIONS FOR FINLAND

On 14 September 2012, the RIPE NCC essentially ran out of IPv4 addresses. They began to allocate IPv4 address space from the last /8 of IPv4 addresses with a very restrictive policy [71]. Before this inevitable end had arrived – it is apparent that – ISPs have rushed into abundantly acquiring IP address blocks.

The BGP-derived data shows that when a prefix is advertised within the routing tables, it does not necessarily match the exact size of an allocated IP address block. There are three ways that a prefix represents an address block: as allocated (exact match), as a fragment of a larger address block, or aggregation of multiple allocated address blocks [72]. In Figure 11, the *matched prefixes* present the IP address blocks that are announced in routing tables as the exact form as they were issued by RIR. For various reasons, operators incline to split up allocated blocks into number of sub-blocks and announce them separately; *fragmented prefixes* present these subdivided blocks.




FIGURE 11 - MATCHED AND FRAGMENTED IP PREFIXES

An announced prefix and IP address block it represents can be a sub-block of another existing (larger) prefix. In such cases, the former is called a *covered prefix* and the latter as *covering prefix*. Covered prefixes are typically used to execute specific load balancing or traffic engineering goals. In Figure 12, the large number of overlapping IP addresses is presented. The 1st level covered addresses are the prefixes in the BGP announcements that are duplicated by exactly one larger prefix. Similarly, the 2nd+ level covered prefixes are the ones that are duplicated by at least two larger covering prefixes.




FIGURE 12 - COVERED IP ADDRESSES

Duplicate announcements of IP blocks are one of the reasons for ever-ascending size of global BGP routing table. As it is apparent in Figure 12, the prefix duplication has been a common practice on the Finnish portion of the Internet. The need for covered prefixes is evident to accommodate various routing preferences. However, as [73] indicates, more than 70% of all BGP announcements captured from public view belong to multihomed stub networks. While the global routing table is often employed to serve local routing interests, it is unlikely that the world outside the national boundaries would follow widely divergent routing paths to reach multihomed stub networks.

The continued growth of the BGP routing table size raises concerns regarding the stability, scalability, management, and increased complexity of BGP operations [74]. In order to keep the routing table growth in check, each AS should announce as few routing prefixes as possible, since each allocated address block will eventually be advertised. Although there is naturally a close correlation between newly allocated IP addresses and BGP routing table growth, other operational factors that contribute to the growth, such as finer fragmentation, multihoming, load balancing should be investigated thoroughly for each Internet region.

Announced block sizes over the last decade provides us with a view about the operational maneuvers of the operators. In Figure 13, announced prefix length distribution is presented. Announcing IP addresses within /24 subnets is by far the most preferred prefix distribution among BGP speakers. Meanwhile, there has been almost no change in blocks larger than /16 prefixes.




FIGURE 13 - PREFIX DISTRIBUTION

On the verge of a shift to IPv6 era, IPv4 address exhaustion has reached a turning-point in 2012 as the stock of previously-unused IPv4 addresses has reached depletion in some regions – the remaining RIRs are expected to deplete their pools soon as well. IPv6 prefixes that are announced from Finnish ASes during the last decade are also captured in our analysis. In Figure 14, we can observe an apparent bump of hype in 2003 and 2004, abiding curve until 2009, and a gradual linear increase for the ensuing years.




FIGURE 14 - IPv6 PREFIXES IN BGP ANNOUNCEMENTS

The transitional technologies – e.g., NAT and dual stack – that will help the Internet to migration from IPv4 to IPv6 require each network to use some IPv4 addresses for backwards-compatibility, as well as the IPv6 addresses. This concerns both new market entrants and previously-established networks. Considering the IPv4 backwards-compatibility requirements of new market entrant network service providers, national regulators may wish to collaborate with their RIRs, and even participate in the IP address allocation policy development process before it causes severe disadvantage for entrants.

6 A Nation-Centric Approach for Studying Internet Interconnections

In this chapter, initially the nation-centric approach for studying the Internet interconnections is introduced. Subsequently, the implementation of this approach to Finland and the results of the longitudinal analysis, along with the classification and inference of relationships between stakeholders of the Finnish Internet industry, are presented. Relevant findings and observations that are obtained with the help of interviews and additional methods to the main BGP derived analysis are also set forth and assessed.

6.1 Deriving a Nation-Centric Internet

Before diving into grappling with voluminous BGP data, a background research about the Internet market in Finland is crucial. As elaborated in Chapter 2, ASes are separate networking and economic entities; technical interaction between each of them naturally reflects economic and organizational level of interactions. Without comprehending the roles of the organizations, ASes are just 32-bit numbers devoid of any content and logical connection to reality. Thus, AS-to-organizational mapping is a constitutive and initial part of this study. Thereafter, the filtration process of the global BGP tables with two complementary methods, and the classification of relevant ASes for a nation-centric Internet are explained.

6.1.1 AS-to-Organization Mapping

On this stage, RIPE IRR database is utilized in order to discover AS-to-organization relationships. A resolute research striving “to develop an organization level view of the Internet’s AS ecosystem”, [75] states that IRR database should be treated cautiously due to the ample amount of incomplete and stagnant data. By cross-checking and hence fortifying the validity of the data with other databases of RIPE – such as LIR list and LIR locator services based on organizations’ address information – IRR data is proven to be valuable for associating relevant ASes with organizations incorporated in Finland. The object names in WHOIS queries including “as-name”,

“descr”, “address”, “role” and “mnt-by” fields are searched through to find out ASes that include any specific information or abbreviations (e.g., FI, Oy and Oyj) related to Finland. Detailed information about queries and object names can be found in [76]. Matching records are extracted and investigated (in the case of matching multiple ASes for an organization) to establish a thorough AS-to-organization matching for relevant ASes.

There are several hurdles of one-to-one mapping of ASes with organizations. An organization may use multiple ASes to employ different routing policies for its internal networks, or it may have acquired several ASes as a result of acquisitions or mergers. Multiple ASes can be also implemented to announce different sets of prefixes at different exit points of each network or to balance traffic across overloaded links. On the other hand, identifying multiple ASes on this stage enables us to infer so-called sibling, s2s relationship between ASes.

6.1.2 Extracting the Data

The process of deriving a portion of the Internet proceeds with filtering the collected BGP data by parsing either AS paths or announced prefixes. There are two approaches applied for filtering the global BGP data that contains millions of lines of BGP messages in order to scale it down to Finland: AS-originated filtering and prefix-originated filtering. Primarily, Python programming language is used for implementing these filtering programs.

Prefix-originated analysis scrutinizes the announced prefixes of each BGP message and searches for IP addresses allocated with FI country code. Filtration works in such a way that it does not only look for exact matching prefixes, as they are allocated in RIRs, but looks for all the possible routable IP addresses from the Finish IP address pool. Having IP address based technique, instead of IP prefix based one, enables us to identify much more announcements that do not include the exact matching allocated prefixes but include somehow readjusted blocks of addresses.

Likewise, in AS-originated analysis, the program assigns a filtering file of our choice, which could be a list obtained from the prefix-originated analysis or a separate list of our choice. Without any elaboration, the filtering program goes over the ASes paths in

each collected BGP announcement to capture the ones containing an AS from the filtering list.

The AS-originated and prefix-originated approaches are used as complementary for each other, especially for the historical view, where the number of databases that reserve historical data is scarce. Discarding past entries, most of the databases provide up-to-date data, which causes to intricate the method used for discovering AS links in the backward-looking analysis. However, we are not in complete darkness; allocated and assigned prefix information is always registered by RIRs with the registration dates [77]. This characteristic makes the RIR delegation data historically retraceable.

The main idea is to take the delegation data as the base and enlarge the AS filtering list in order to capture the ASes that used to be active but presently inactive. In the longitudinal analysis, each year yields different set of list for the previous year's analysis. For the sake of comparing annual changes in the lists pertaining the sequent years, revising each year's list and carrying on with the filtration process accordingly are vital. Withal, as long as observed results are not verified with the historical and conventional facts of the industry, findings of the analysis would remain devoid of cogency. The flow chart that is used in the longitudinal analysis for expanding the filtering list of ASes can be found in Appendix A.

6.1.3 Stakeholder Categorization and Clustering

The organization mapping endeavor and realizing how nebulously entwined the boundaries of nation-state regions of the Internet become make it apparent that actualizing a notion to designate a single and all-inclusive list of ASes that share identical properties in a country is a futile struggle. Therefore, the method used in this study proposes a scheme that situates ASes that are relevant for the Finnish Internet in three interconnected clusters.

Firstly, we can start with listing stub networks and ASes that only announce prefixes with FI code since this listing can be conducted in a relatively straightforward manner. When considering the routes in Internet to reach destinations in Finland, the rightmost ASes in AS paths are where the announced prefixes belong to. These ASes are where we can find stub networks – i.e., the edges of the Internet – that do not

provide transit for other ASes. Thus, if an AS appears only as the rightmost AS in the collected paths, then it can be classified as a stub network. The organizations that administrate those ASes have to be incorporated in Finland. In no circumstances shall the networks of these organizations be presented abroad; for instance, if they engage in public peering in IXPs, they shall only do so in Finland. Typically, these are Finnish ASes, and undoubtedly, they announce only Finnish prefixes allocated with FI country code. Hence, they will be referred to as *the first level relevant Finnish ASes* throughout the thesis.

Secondly, there are multinational ASes which announce prefixes allocated from both Finland and other countries. These ASes may publicly peer outside of Finland and provide transit for the first group of ASes. Mainly, these *second level relevant Finnish ASes* belong to the players the roles of which are pivotal for the Finnish Internet market – unlike the deception that the degree notation in their designated name may cause. When solely mentioned *Finnish ASes* without any relevance level, then it refers to the first and the second level relevant Finnish AS lists, hereafter.

Lastly, there are *the third level relevant ASes* that are transit providers for either one of the previous two AS clusters. These ASes are at the core of the Internet and some of them might be even global Tier-1 ASes, explained in the Chapter 2. Their networks may or not have a point of presence in Finland. However, since they are relatively bigger players in the global Internet transit market, intrinsically, they do not engage in peering relationship with any BGP speaker from Finland. It is, to a certain extent, difficult to define the complete list of these ASes; therefore, the discovery of this list of ASes is left out of the scope of the longitudinal analysis.

6.2 AS Links: Business Relationships

Comprehending the ecosystem construct also requires inferring the type and function of linkages between ASes. The procedure for a novel and simplistic business relationship inference method for observed peering and transit links in a national part of the Internet can be broken down to steps as follows:

- Extract relevant AS adjacencies from BGP data,
- Purify each AS path from duplicates,

- Apply AS path-based inference analysis,
- Apply AS policy-based analysis and utilize the additional information.

From those AS paths with relevant adjacencies, all unique AS links business relationships of which are aimed to be inferred are extracted. The eclectic approach comprises several parts; succinctly it is based on utilizing two BGP attributes: AS path and AS policy information such as BGP communities and local preferences.

6.2.1 The Inference Approach

Due to the fundamental deficiencies of the collected BGP data, it is neither feasible nor possible to infer all linkages between ASes. Studies, such as [11] and [78], on the completeness of observed AS-level Internet indicate that the percentage of missing links in the BGP-derived data – which are mostly peering relationships – can range from 10- 20% for Tier-1 and go up to 85% for Tier-2 ASes. Moreover, BGP data may miss up to 90% of peering links in the case of large content provider networks, which have been increasingly establishing new peering links in the recent years.

The same studies also suggest that through the use of data collected over long enough time, collected BGP data can capture all the customer-provider (transit) links in all tiers of the Internet. The underlying rationale of this assertion relies on the fact that the BGP collecting projects have monitors in all the Tier-1 ASes except one, however, that particular Tier-1 AS has a customer that hosts a monitor. By definition, Tier-1s only peer with each other in a mesh topology and do not need to buy transit from any other AS to reach any destination. Hence, we should be able to observe the complete list of all transit relationships.

Thus, within a confined scope, the objective of our inference method is to discover any possible peering relationships that might be captured from the BGP data – especially the ones belong to the monitor ASes (BGP announcement originators) or to upstream providers of those monitor ASes. After inferring the peering relationships, accordingly, transit links are defined to depict inter-AS relationships. In other words, we subtract incompleteness from the rest of the data to reach completeness.

6.2.2 The Properties of AS Path

Initially, it is beneficial to investigate the properties and dynamics of AS path in BGP messages. The AS-path-based analyses lead to interpretations on the relationships that heavily rely on probability theory – probabilities of position and frequency of AS links and individual ASes appearing in unique paths.

In an AS path, the leftmost AS is always the originator of the announcement where the monitor router is located and the route collector is provided with the routing tables. The rightmost AS of the path is the destination AS where the prefix is announced from. ASes usually receive path information to the same prefix from multiple neighboring ASes. In general, ASes tend to prefer the path advertised by a customer (p2c) over that from a peer (p2p), and similarly prefers a path from a peer over that from a provider (c2p). This is referred to as no-valley and prefer-customer policy [58], as previously presented in Section 4.3.1. Even though this policy is not a rule that all ISPs obey, it is a very common practice among the participants of the Internet. Yet, there are studies that argue for the opposite. As stated in [79] attempting to quantitatively characterize BGP announcements that violate this valley-free property, “the valley announcements are more pervasive than expected”. Professedly, a considerable number of examples for intentional valley announcements are uncovered. Nonetheless, this property of AS interconnections has predominantly been taken for granted in inter-AS business relationship inference studies.

Figure 15 presents the possible relationships for each link of an AS path comprising of the following sequence: “AS1 AS2 AS3 AS4”. Mainly, the position and the order of these ASes are stressed in the illustration. The link type “c2p” with inclined slope represents customer-to-provider relationship. Peering relationship is represented with “p2p”, whereas “p2c” stands for provider-to-customer relationship. This particular AS path consisting of four ASes may yield seven distinct scenarios wherein seven different sets of business relationships exist. Among these seven scenarios, we may infer that between AS3 and AS4 it is very unlikely to observe a peering relationship – still, dependent on previous ASes in the path. On the other hand, if AS1-AS2 link only appears on only one of the first steps of the path and disappears later, we may infer that the link can be a peering link. Succinctly, the further positions from AS paths are extracted, the less likely to encounter peering relationships.




FIGURE 15 - REPRESENTATION OF POSSIBLE BUSINESS RELATIONSHIPS IN AN AS PATH COMPRISING OF FOUR ASES

However, we cannot solely rely on positional probability; there are many factors that alter the observed routes. For example, in order to capture a BGP messages from the public view that indicates a peering link, one of the peering ASes of which has to either be monitor AS, or have a downstream customer that provides the collectors with the routing announcement. Conspicuously, this condition cannot be catered for all ASes due to the no-valley routing policy and the impracticability of deploying monitor routers in stub networks. This is the main reason that it is not plausible to infer the complete set of peering relationships from the public view.

Moreover, the RIS route collectors receive both full routing tables and partial views, depending on monitor ASes. For example, if there is a multi-hop BGP session with a monitor AS, a broad view of the routes is seen at only one location while providing a partial view of routes on other collectors – in order to prevent unnecessary overload on RIB sizes of collectors. After all, a monitor AS may choose to not share some of its routing table information arbitrarily or due to non-disclosure requirements [45].

6.2.3 Observed AS Relationships

Presently, RIS project has approximately 400 monitor ASes peering with its RRCs, out of over 40,000 existing ASes globally. This ratio gives a rough idea about the percentage of peer links that might be missing from the public view. For the case of Finland, there are only 14 monitor ASes from all three lists of relevant ASes – 4 of them from the second level relevant group and 10 of them from the third level relevant group.

With the help of filtration methods, only relevant and unique AS paths from each RRC are extracted in such a manner that it will not be necessary to be inundated with avalanche of data. For instance, at RRC1 collector located in London Internet Exchange (LINX), currently a RIB – updated routing tables in every 8 hours for each collector — comprises of more than 5 million lines of BGP messages in total. However, the number of BGP messages that is needed to study the Finnish AS ecosystem can be boiled down to about 22,000 (including duplicates). Overall from 17 RRCs, approx. 18,000 unique AS paths are harnessed from all monitors to reveal 326 unique links between the Finnish ASes.

After extracting unique AS paths from collectors, each AS path is cleared of sequential repetitions. The reason for these repetitions is a traffic engineering method called AS-prepend. Operators tend to extend the length of AS path by defining next hop as their AS number so that their providers would not choose to have the traffic through their networks due to shortest path preferences [80].

On this phase, an inference table, the partial snapshot of which is presented in Table 2, is generated. The peering detection algorithm is conducted through considering several parameters in a panoptic manner. In repetition (*Rep.*) parameter, the number of repetitions of each AS link in AS paths is analyzed while evaluating uniqueness of these AS paths with the unique path (*Unq_Path*) parameter. The sequencing parameter – in which part of the sequence and how many times AS pairs (AS1 and AS2) appear – is presented within *AS path sequence* column. In *AS1_attribute*, same sequencing and repetition analyses are conducted for AS1, which is the left AS of the adjacency. This attribute is crucial to evaluate the weightiness of the particular AS and provides specific indications to interpret if a monitor AS has also a downstream

monitor AS, which increases the likelihood of encountering a peering relationship on the further steps of the AS path.

Additional to the AS path information, there are BGP community and location information that are made use of during the analysis. Location information is an important supplement for discovering peering relationships; 35 ASes that have participated in Finnish Communication and Internet Exchange (FICIX) during the last decade are listed and marked. Hence, we could also pay special attention to the links between those ASes presences of which reciprocally coincide with each other's in FICIX.

BGP communities attribute may also include location, local preference and type of interconnection information within non-uniformly coded forms for each ISP. BGP communities attribute is an intact part of the BGP attributes, which have not been systematically exploited for relationship inference. The lack of taxonomy and standardization in BGP communities, and the voluntary (sometimes haphazardly) usage of these attributes have been the major reasons for this neglect. Usually, operators share the usage and meaning of their communities on their IRR entries or on their websites. Even if they are not explicitly expressed, it is still possible to extrapolate the meaning of a pertinent link from other observed links. In the up-to-date view of the analysis, there are 326 observed links that interconnect Finnish ASes to each other and to the rest of the Internet. 200 of those links (61%.3) have provided BGP communities attributes that help us observe the type of business relationship

AS1_Attribute		AS1	AS2	Rep.	Unq Path	AS PATH SEQUENCE					Community	MED	RRC	Inference
						1st	2nd	3rd	4th	5th				
Monitor	1st	8359	47605	140	37	35	2	0	0	0	8359:38359:5002	0	rec12	Peer detected
Monitor	2nd	9002	29422	156	38	13	24	1	0	0	9002:9002:9002:64623	0	rec07	Peer detected
Monitor	2nd	6667	1759	347	36	36	0	0	0	0	6667:3001:6667:4002:6667:5201	20	rec07	Peer detected
Monitor	2nd	9002	719	853	59	20	39	0	0	0	9002:9002:9002:64623	0	rec07	Peer detected
Monitor	3rd	16086	42343	7	4	1	1	2	0	0	#N/A	#N/A	rec00	
	3rd	719	197930	47	41	0	1	26	10	2	719:900:719:1000:719:2000:719:3000:719:4002	0	rec00	
Monitor	2nd	3292	51311	47	41	1	22	15	3	0	3292:1000:3292:23500	0	rec00	
	2nd	174	198514	47	44	0	28	12	3	1	174:1001:174:21101:174:22014	0	rec00	
	4th	1759	30798	48	28	0	4	9	12	2	1759:100:1759:103:1759:30010	20	rec00	

TABLE 2 - SAMPLE FROM INFERENCE SCREEN

more thoroughly. For example, Elisa's AS 6667 defines a community value "6667:3000" corresponding to the routes received from its customers, whereas the routes to its peering community are defined with "6667:3001".

6.3 Results of the Longitudinal Analysis

For the latest view of the AS-level interconnections, the prefix originated filtering of the BGP public view according to 2012 delegation data yields 192 ASes that are announced as the originators of the Finnish prefixes. After applying the explained organization mapping methods with the help of the supplementary databases, we can define 167 of these ASes as the first level relevant and 9 of them as the second level relevant Finnish ASes. As an attempt to observe and understand the evolution of the Finnish Internet ecosystem during the last 10 years (2002–2012), the methods are abided by and carried out for a retroactively longitudinal analysis. In Figure 16, the number of observed ASes that are actively used in each year is presented.

At this point, a new dataset enriched with additional attributes of ASes is generated for each year. These attributes are the names of the organizations with a sectoral categorization, number of routable IP addresses announced from each AS in order to estimate the size of the networks that ASes administer and AS adjacencies to provide the necessary data for each year's inference tables, as explained with Table 2. Incorporating diverse internal properties should become an indispensable component of nascent AS studies which are committed to strive towards constructing viable and realistic AS graphs, not featureless clusters of identical nodes.




FIGURE 16 - NUMBER OF ACTIVE ASes

The categorization of organizations that administer their own public ASes is a significant part of our attempt to provide internal specifications of each AS. The Global Industry Classification Standard (GICS) structure, consisting of 10 sectors, 24 industry groups, 68 industries and 154 sub-industries, is used for categorization [81]. The classification standard unobtrusively provides both compactness and granularity that is aimed to draw on.

After combining the findings that are obtained for the latest view, the analysis reveals that there are 165 domestic and international entities that administer above-mentioned first and second level relevant list of ASes, and engage in BGP interconnectivity in Finland. A pie chart of the all 165 organizations divided into sectors is shown in Figure 17. The diversity of the stakeholders in the ecosystem is remarkable – from universities, to governmental organizations, major telecommunication networks to enterprises with varying industrial backgrounds.

The significant take-away from this categorization is that more than one-fourth of the organizations that manage their own ASes are the companies that are outside of the ICT sector. Even though there has not been any study on sectorial categorization of the Internet management, the ratio is higher than one could expect.




FIGURE 17 - STAKEHOLDERS OF BGP INTERCONNECTIVITY CATEGORIZED BY SECTORS

6.3.1 Inference of Observed Links

Inference method of AS adjacencies and business relationships leads us to a number of transit relationships between Finnish ASes. Exposing the customer-provider relationships between these organizations indicates the direction that the money flows in ICT sector for Internet transit agreements. For visualization of our dataset, Cytoscape [82] – an open source platform for complex network analysis and visualization – is used. Figure 18 represents the latest view of the AS clusters with their transit interfaces in a circular layout – the inner circle of ASes constitutes the 3rd level relevant cluster while the stub networks are incorporated in the outer circle. A hierarchical layout of the transit relationships can be found in Appendix B.




FIGURE 18 - CIRCULAR LAYOUT OF AS CLUSTERS

As shown in Figure 19, a complete set of transit business relationships are revealed for Finnish ASes, whereas the list of peering relationships that is observed from the public view BGP data is far from being complete. Nevertheless, the list of observed peerings is employed later while generating a public peering matrix, with several other methods, explained in Section 6.3.5.




FIGURE 19 - OBSERVED AS LINKS

6.3.2 Multihoming

Most of the large-scale enterprises spread their Internet traffic across two or more ISPs in order to improve performance and resiliency of their networks, and reduce costs in a competitive market. In Figure 20, we can observe that purchasing Internet transit interconnectivity from more than one upstream provider to deliver and receive a portion of Internet traffic has been prevalent among Finnish ASes.

In the Finnish Internet ecosystem, multihoming has definitely not been an evanescent practice during the last decade. The practices of stub networks deploying BGP multihoming, which constitutes merely a limited portion of all multihoming practices (regarding NAT multihoming), prove this claim. However, the interviews conducted during the research reveal that for stub networks, there are some hurdles that they need to overcome. Many of these enterprises do not have the internal networking resources to exploit the advantages of multihoming, such as route optimization, and reliability and performance improvement, elaborated in [83]. As Internet transit prices continue to decline, nowadays operational costs including training and employing proficient networking staff for 24/7 network operations have surpassed the wholesale Internet transit prices and have become one of the leading cost factors.




FIGURE 20 - AVERAGE NUMBER OF UPSTREAM PROVIDERS

6.3.3 Market Trends

The most significant insights we grasped from the interviewees are on the aspects related to pricing and agreement issues of interconnections that could not be studied by the technical means that are provided by the open nature of the Internet. The pervasive view indicates that the transit pricing have become so low that transit costs do not constitute the main cost driver anymore. For most of the cases, buying transit has started to be more sensible unless there is a set of customers that need to be reached directly from a local AS. Therefore, ISPs have started to reconsider their motives for previously established peering and transit agreements. As a coincidence, there have been several de-peering and paid peering agreements taking place, associated with the competitive pricing and changing market conditions. Especially, large and medium-sized ASes have started to pursue more selective peering policies when choosing their peers in order to avoid establishing peering relationships with potential transit customers.

For each year, the top five transit providers are ranked by the number of their downstream customers in Figure 21. Besides the tremendous increase in the number of transit agreements signed by top five providers, we can observe the apparent reflections of organizational changes in the interconnectivity market on the AS-level. However, morphings in the inter-domain world are subject to a more protracted process. Although the companies may rise, shine and fade away from the consumer market relatively easily, AS names and the customer list of ASes that once belonged to a merged or acquired company are resumed for ensuing years. Positional changes in the ranking of these top ASes are not accomplished over a short period of time; rather, operators often take up with the heritage of their predecessors.




FIGURE 21 - TOP FIVE TRANSIT PROVIDERS OF THE FINNISH INTERNET TRANSIT MARKET

6.3.4 Public Peering in Finland

Finnish Communication and Internet Exchange association is the biggest Internet eXchange Point (IXP) in Finland. Since 1993, FICIX has been functioning as a non-profit organization, in which neutral and reliable peering facilities for its members are provided. Currently with 28 members, FICIX operates at three different locations – FICIX 1 at City of Espoo, FICIX 2 at City of Helsinki, FICIX 3 at City of Oulu – which are not interconnected to each other [84].

FICIX is an apt representative of the European model for IXPs, which is largely based on the successful model initiated by the LINX. In the “classic LINX model”, the co-location provider may subsidize or pay for elements of having the IXP within their facilities (e.g. space, power, fiber, equipment costs). The bottom line of the model relies on the separation of peering fabric operator and the co-location operator – i.e., keeping the “co-location neutrality” where customers can choose any co-location facility that meets their requirements when peering at IXPs.

On the other hand, Tampere Region Exchange (TREX) is an academy-boostered, next generation Internet exchange point in Tampere, Finland [85]. Even though the term “the next generation” is a marketing hype, it is an all-encompassing term per se to describe the means of value proposition to acquire more members/customers to join the IXP. TREX strives to differentiate from traditional IXPs model, by allowing, and to some degree, by even endorsing its members to do business together (e.g., buying bit streaming services or transit, setting up private VLANs).

FICIX and TREX are both members of Euro-IX, which gathers up to 60 IXPs from 36 countries worldwide [55]. Traditionally, the Finnish Internet traffic has nearly doubled every year, and the level of Internet traffic per capita is relatively high, compared to the central European Internet traffic [84]. IXPs in Finland further expand the interconnectivity and robustness of the Internet as a whole. Considering the average traffic carried over IXPs in Finland in 2012, the distribution of the traffic load is presented in Figure 22, which should give us a rough estimation about the size and the functionality of each IXP.




FIGURE 22 - PUBLICLY EXCHANGE IP TRAFFIC IN FINLAND

The global IXP ecosystem has also become tiered, presenting an analogy to the hierarchical model of the relationships between ASes. There are a handful of Tier-1 IXPs (e.g. AMS-IX, DE-CIX and LINX) where one can peer with ASes from all over the globe and one can realistically expect to get one or two hundred thousand routes by peering. Then, there are Tier-2 IXPs (e.g. Netnod, NIX, and MIX-IT) where one can establish peering with international networks or ISPs from neighboring countries. For example, Netnod in Stockholm connects all Nordic ISPs together and in addition several Russian, German, and UK based operators. Finally, there are Tier-3 IXPs (e.g. TIX in Tallinn and Netnod's exchanges outside of Stockholm) where one can only hope to exchange traffic within a very limited scope. FICIX 3 and TREX obviously

are Tier-3, and the main reason is location. The foreign ISPs are only willing to set up gear in densely populated cities, i.e., at the peering hubs.

Emphasizing the importance of location and population density and validating the traffic volumes of public peering points in Finland, an OECD report [12], depicts the gradual progression of the largest European IXPs towards the center of the population density. The argument advocated in the report indicates that average distance between an exchange point and the population that its members serve is a major measure of the efficiency and the long-term success of any IXP. Therefore, all other factors aside, exchanges nearest to the centers of population density tend to be the most successful. The amount of exchange traffic that they generate can be carried to their constituents at the lowest possible cost while retaining the highest possible performance.

Naturally, Amsterdam gradually overtook London, and Frankfurt gradually overtook Amsterdam. For Finland, it is apparent that a similar correlation between IXP locations and population density applies. Because of the multi-lateral peering agreement⁶ history, which has been abandoned for almost a decade, Finland has had a well-propagated peering exchange matrix for domestic routes. However, the same historical reasons have held FICIX 1 and FICIX 2 back from becoming a Tier-2 IXP. After the rules of the national exchange point have become more flexible, there has been an influx of international market players joining these exchange points, which is nowadays promoting FICIX 1 and FICIX 2 towards the higher tiers of the global IXP hierarchy.

6.3.5 Peering matrix in FICIX

As [11] indicates “the percentage of missing peering links in the BGP-based data can range from 10- 20% for Tier-1 and Tier-2 ASes to 85% or even more for large content networks”. However, the same study also suggests that through the use of BGP data collected over long enough time periods, “the public view can capture almost all the customer-provider links at all tiers in the Internet”.

⁶ Multi-lateral peering agreement is an agreement whereby all signatories agree to peer with all other signatory parties of the same agreement.

On the other hand, another study reveals that even a laborious attempt to collect publicly available BGP data with hard-to-get non-public control-plane measurements and the latest available state-of-the-art data-plane measurements could only account for revealing 30% of all known peering links at a large European IXP [86]. The authors deduce that the available historic datasets and current methods are not sufficient to infer the evolution of peering links.

Another significant research on IXPs, [87] aims to provide a more comprehensive and complete picture of the IXP substrate with combining several different methods. In fact, their results provide a detailed account of the information and efforts needed to discover and map each and every IXP worldwide. Their exclusive focus on IXP consists of launching targeted trace-routes via looking glass servers ⁷ from meticulously selected sources to chosen destinations and accordingly checking the resulting paths for indications whether packets go through an IXP. By combining our BGP-derived inference and IRR data with the same methods proposed in [87] for FICIX, 149 unique peering relationships are revealed in the FICIX peering matrix which is significantly greater than the number that any study has been able to achieve so far.

The observed FICIX peering links are depicted within a matrix format in Figure 23. The more sources of conformation for a particular peering relationship between AS pairs are obtained, the darker blue the color of the particular cell gets. Orange colored cells indicate the occurrence of a transit relationship between AS pairs.

⁷ In FICIX, 4 out of 28 members provide looking glass servers, which are publicly accessible servers for performing routing queries used to troubleshoot routing issues across the Internet.

6. A NATION-CENTRIC APPROACH FOR STUDYING INTERNET INTERCONNECTIONS

FICIX	702 Verizon	719 Elisa	1248 Nokia	1342 Fujitsu	1741 Funet	1759 TSF	2686 AT&T	3292 TDC	5400 BT	6667 EUnet	9002 RETN	12552 IP-Only	16086 DNA	20569 Alnacom	24751 Multi.fi	29154 Academia	29422 Nebula	30798 TNNet	34188 TeleKarelia	34384 WLANnet	47605 FNE	8359 MTS	33926 EuroTr	12659 EuroBNN	43646 Smartlog
702 Verizon										IRL BGP	BGP		IRL									BGP			
719 Elisa			IRL, I	IRL	IRL	IRL		IRL, TR		T	BGP		BGP					T							
1248 Nokia		I		I	IRL, I	T		IRL, TR		T	I	I	I	I	I	I	IRL	I	I	I	I	I		I	
1342 Fujitsu		IRL	I		IRL	IRL	IRL	IRL, TR		T	BGP			IRL			IRL								
1741 Funet		IRL	IRL, I	IRL				IRL, TR	IRL	IRL BGP	IRL BGP	IRL	IRL BGP	IRL	IRL	IRL	IRL	IRL	IRL		IRL	IRL BGP		IRL	IRL
1759 TSF		IRL	T	IRL	IRL			IRL		IRL BGP			IRL BGP	IRL	IRL	IRL	T	T							
2686 AT&T	IRL			IRL				IRL		IRL, TR, BGP	BGP				IRL	IRL	IRL					BGP			
3292 TDC	IRL, TR	IRL, TR	IRL, TR, I	IRL, TR	IRL, TR	IRL, TR	IRL		IRL, TR	IRL, TR	IRL, TR, BGP	IRL, TR	IRL, TR	T	IRL, TR	IRL, TR	IRL, TR	IRL, TR	IRL, TR		T, I	IRL, TR		IRL, TR	
5400 BT					IRL			IRL, TR		IRL, TR															
6667 EUnet	BGP, IRL, TR	T	IRL, TR, I	T	BGP, IRL, TR	BGP, IRL, TR	IRL, TR, BGP	IRL, TR	IRL, TR		IRL, TR	IRL, TR	IRL, TR		T	T						IRL, TR	IRL, TR		IRL, TR
9002 RETN	BGP	BGP	I	BGP	BGP		BGP	BGP, TR		IRL, TR			BGP	BGP	BGP	BGP	BGP	BGP	IRL	T	T				T
12552 IP-Only			I		IRL			IRL, TR		IRL, TR				T		IRL	IRL		IRL						IRL
16086 DNA	IRL	BGP	I		BGP	BGP		IRL, TR		IRL, TR	BGP			BGP		BGP	BGP				BGP	BGP		IRL	
20569 Alnacom			I	IRL	IRL	IRL		T		IRL BGP	T?	IRL BGP			IRL	IRL	IRL	IRL	IRL	IRL	IRL	BGP		IRL	IRL
24751 Multi.fi			I	IRL	IRL	IRL	IRL	IRL, TR		T	BGP			IRL			IRL	IRL	IRL	IRL		BGP		IRL	
29154 Academia			IRL, I	IRL	IRL	IRL	IRL	IRL, TR		T	IRL BGP	IRL	IRL BGP	IRL	IRL		IRL	IRL	IRL	IRL	IRL	IRL BGP		IRL	IRL
29422 Nebula			IRL, I	IRL	IRL	T	IRL	IRL, TR			IRL BGP	IRL	IRL BGP	IRL	IRL	IRL		IRL	IRL	IRL	IRL	IRL BGP			
30798 TNNet		T, I	IRL, I		IRL	T		IRL, TR			IRL BGP			IRL	IRL	IRL	IRL	I	IRL	IRL	IRL			IRL	
34188 TeleKarelia		IRL	IRL, I		IRL			IRL, TR			IRL	IRL		IRL	IRL	IRL	IRL				T	BGP		IRL	
34384 WLANnet			I							T											T			IRL	
47605 FNE			I		IRL			T, I		T			BGP	IRL		IRL	IRL	IRL	T	T		BGP		IRL	
8359 MTS	BGP		I		BGP		BGP	IRL, TR		IRL, TR			BGP	BGP	BGP	BGP	BGP		BGP		BGP				
33926 EuroTr									IRL, TR																
12659 EuroBNN			IRL, I	IRL	IRL			IRL			T	IRL	IRL	IRL	IRL		IRL	IRL	IRL	IRL	IRL				
43646 Smartlog					IRL				IRL, TR					IRL		IRL									

FIGURE 23 - FICIX PEERING MATRIX

7 MNO Interconnections in Finland

With the migration to IP-based inter-operator interconnections, MNOs and international carriers are facing new opportunities for interconnection architectures, rather than simply replacing old TDM interconnections with IP pipes. As IP-based networks evolve, different models of interconnection between operators are being shaped both at national and international levels. Activities on the international level are developing in a generally more organized fashion as discussed in Chapter 3. At the national level, regulators are actively taking the leading role, defining the future interconnection landscapes in their own milieus. In cooperation with national operators, Finnish Communications Regulatory Authority (FICORA) has also undertaken a responsibility for establishing an IP interconnection exchange in Finland.

As an initial attempt towards this objective, in 2006, FICORA began to register Finnish public ENUM (E.164 Number Mapping) domains, founding process of which has been recognized by several telecom industry reports on VoIP/ENUM (e.g., [88]). By definition, public ENUM enables end users and companies to control their own means of communication when necessary and, for instance, receive calls dialed with telephone numbers even directly over the Internet. This form of ENUM was thought to cut out the role of operators from the process, and bring cost savings and more versatile VoIP call features to end users. However, the Public ENUM registries remain active in only 12 countries [89] and have not reached the projected success. The main reasons for that would be the lack of available services and commercial interest.

In the following years, a SIP working group was established by FICORA with collaboration from the leading operators in Finland. The main goal of the working group was to analyze the feasibility of a (private) operator ENUM project along with a SIP/VoIP interconnection project which is commercially guided by Finnish number portability management company, Numpac. As an outcome of these endeavors, Domestic IP eXchange (DIX) is established between telecom operators to enable IP interconnection for voice traffic.

As shown in Figure 24, the operator networks essentially provide a private IP interconnection between each other, which requires bilateral agreements for settlement between each party. In this pure IP peering architecture, operator networks are agnostic to the type of application or media exchanged through the network. Thus, operators typically apply charges based upon the amount of data transferred within a traditional CPNP charging model.

The IP interconnection architecture endows the network operators with flexibility in choosing the most efficient path for particular destinations according to the packet traffic and commercial requirements. In DIX architecture, there are two points of interconnection, which are not connected to each other in order to prevent routing loops, and a third one is also envisaged to be built. From an economy of scale standpoint, the benefits in terms of reduced administration and interconnection fees can be achieved. TDM interconnections require much higher number of point of interconnection established between operators, yet provide less effective results in terms of data exchanged. In order to incorporate technical expertise developed at the international level and due to interoperability concerns, national IP interconnection practices more or less technically mirror international interconnection approaches [90]. Similarly, in DIX, the defining rules for interconnection requirements and maintenance rules cohere with the GRX/IPX guidance of the GSMA.




FIGURE 24 - DIX INFRASTRUCTURE

Open to all network operators that meet the requirements defined by GSMA recommendations, such as [30] and [91], DIX interconnect currently provides service for only voice over IP packets. Once IP-based interconnection for voice is established, it paves the way for a natural evolution from voice to multimedia interconnections across MNOs, providing end user services such as instant messaging, presence and video sharing, i.e. RCSs, as discussed in Chapter 3. The interoperability and assured end-to-end QoS service delivery models between MNOs come into prominence, particularly for the delivery of these end-to-end multimedia services.

With the guidance of clear frameworks to adhere, interconnection architectures (either on national or international level) provide well-defined business and technical roles for MNOs and the global carriers. The argument has been uttered many times by interconnection experts that *off the Internet* type of models are necessary for a sustainable IP interconnection solution to guarantee QoS for enhanced end-to-end multimedia services. Presently, it seems “unlikely that those models based upon interconnections over the public Internet can provide such sustainable solutions” [22].

Today, the sheer scale of the Internet makes it not a trivial task to implement the same QoS mechanisms in every single Internet device around the world. Quality concerns also pose a risk that the present open Internet may evolve into two separate internets. The classical best-effort Internet may remain as it is now where operators limit their investments just enough to maintain end-to-end interconnectivity, and thus preserve the ubiquity of the service, with possibly some minimal service quality guaranteed by the regulations. Meanwhile, within a service aware and private interconnection model such as DIX, premium interconnections can be offered. Operators may presumably invest in the establishment of these new private and secure interconnections that can support QoS assurance while recovering their investments by charging for relatively higher prices.

8 Conclusions

Having looked into the ways to study Internet interconnections in a nation-centric approach, we have considered the issues that are pertaining to technical, economic and regulatory aspects. Implications that the emerging IP interconnections pose for the future global service delivery among both fixed and mobile Internet are argued from the perspective of the existing practices. Within this chapter, exploitation of the results and limitations of the study are discussed, key findings are highlighted, and finally topics for future research are suggested.

8.1 Assessment of Results

The nation-centric analysis, conducted for Finland by using the RIPE data and the RIS project that collects BGP routing data, can be applied for other countries as well. The level of generalizability of this analysis for other countries is highly associated with the level of manual interpretations required due to the non-uniformities in the datasets. For Finland – where only 1.8% of all allocated IP addresses in RIPE zone are used – this can be practicable. However, for countries that have a larger impact on the global Internet (e.g., Russia, Germany and the UK), the framework can be laborious to execute without the means of automatized collection and inference mechanisms.

Prior to the analysis, the current measurability deficiencies on the collection and inference methods of the BGP data are addressed. Accordingly, enhancing the reliability of the raw data is proposed by jointly utilizing the Internet registry data and BGP routing data derived from the actual state of the Internet. Realizing the futility in designating a single list of ASes that share identical properties within borders of a country, a categorization technique of national ASes is suggested within three clusters – in terms of geographical distribution of organizations and announced IP addresses. Consequently, consistent results for the longitudinal analysis are able to be obtained, and the inference on the business relationships is carried out to discover a complete list of the Internet transit arrangements.

The main limitation for the results, as for any BGP-based Internet study, which may jeopardize the integrity of our approach is the deficit in discovering peering links – especially the ones belong to content providers. This limitation emphasizes once again the public data provided for interconnection studies and agreement models are scarce – particularly, for peering and MNO interconnections it is beyond scarce.

8.2 Exploitation of Results

The Finnish Internet market has evolved through many phases since its inception. As we are able to observe them in the longitudinal analysis, during those transitions, mergers and acquisitions have been in the market as natural outcomes of dense competition. The number of gateway ASes that interconnects Finland with the rest of the global Internet has remained stagnant, while the overall number of Finnish ASes has been increasing throughout the past decade. The domestically well-established public peering fabric and the multihomed nature of the Internet transit business have consolidated the interconnectivity market. The operators possessing their own national backbone networks are mostly competing with each other in this market while being challenged by international players who pursue assertive pricing strategies.

For stub networks – the ASes that are located on the far edges of the Internet construct – an important finding comes to light with the observation that multihomed transit agreements have not been an evanescent effort and have been gradually increasing over the last years. Although the operational costs are setting up a high barrier for many Finnish stub networks to exploit multihoming benefits, the practice of multihoming has been a crucial factor to bolster the market competitiveness.

The BGP analysis and our nation-centric approach incorporate the investigation of the interconnections between industry players. Nevertheless, a BGP interconnectivity analysis should not be treated as interchangeable with an Internet traffic measurement analysis by which usage dynamics and traffic volumes are harnessed. The longitudinal BGP interconnectivity study, depicting the evolution of the AS-level ecosystem, should be relevant to practitioners and authorities who desire to comprehend both technical and business interfaces between stakeholders within a national Internet milieu and accordingly design new interconnection policies.

8.3 Future Research

Improved transparency into the mechanisms of the Internet operations, along with its possible contribution on publicly debated and healthier discussions of interconnection related issues, is worthy of further studies. More public data may help the development of frameworks, on either global or national scales, leading to commonly accepted set of assumptions for the outcome of negotiations and inducing a possible reduction of the transactional costs.

Having quality and security concerns prevail over the cost and administrative practicality incentives, IP traffic exchange between MNOs follows a separate, “off the Internet” trajectory. For the emerging interconnection models facilitating the delivery of quality assured end-to-end services, the need for coherent charging mechanisms to correlate retail pricing with wholesale interconnection fees and accommodate equitable distribution of network costs remains a compelling topic for future research.

While MNO models of service aware interconnections that are able to incorporate traffic from ISPs as well are coming forth, more of a philosophical debate, rather than pragmatic, arises. Whether having two separate worlds of interconnection for global IP services delivery is an appropriate and efficient way to substantially enhance the social welfare? A worthwhile research pursuit yielding an answer to this question could help identify the role of interconnection models in terms of both effective and efficient service delivery.

References

- [1] H. R. Varian and C. Shapiro, *Information Rules: A Strategic Guide to the Network Economy*, Cambridge: Harvard Business School Press, 1999.
- [2] Odlyzko, A. and B. Tilly, "A Refutation of Metcalfe's Law and a Better Estimate for the Value of Networks and Network Interconnections," no. March, 2005.
- [3] W. B. Norton, *The Internet Peering Playbook: Connecting to the Core of the Internet*, DrPeering Press, 2011.
- [4] CIDR Report, [Online]. Available: <http://www.cidr-report.org/as2.0/>. [Accessed 3 January 2013].
- [5] P. Faratin, P. Gilmore, S. Bauer, A. Berger and W. Lehr, "Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy," in *Telecommunications Policy Research Conference*, 2007.
- [6] K. Butler, T. Farley, P. McDaniel and J. Rexford, "A Survey of BGP Security Issues and Solutions," in *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100-122, January 2010.
- [7] P. I. Beijnum and I. V. Beijnum, "Interdomain Routing History," in *BGP: Building Reliable Networks with BGP*, O'Reilly, 2002, pp. 18-19.
- [8] Y. Rekhter and T. Li., "A Border Gateway Protocol 4," RFC 1771, 1995.
- [9] R. Oliveira, "Understanding the Internet As-Level Structure," Ph.D. Dissertation, University of California, Los Angeles, 2009.
- [10] APNIC, "Autonomous System Numbers," [Online]. Available: <http://www.apnic.net/services/services-apnic-provides/helpdesk/faqs/asn-faqs#UsePrivateASN>. [Accessed 2 November 2012].
- [11] R. Oliveira and W. Willinger, "The (In)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Transactions on Networkin*, vol. 18, no. 2, pp. 109-122, 2010.
- [12] D. Weller and B. Woodcock, "Internet Traffic Exchange: Market Developments and Policy Challenges," OECD Report, 2012.
- [13] K. Claffy, "Workshop on Internet Economics (WIE2011) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 42, no. 2, pp. 110-114, April 2012.
- [14] D. Clark, W. Lehr and S. Bauer, "Interconnection in the Internet: The Policy Challenge," in *TPRC*, 2011.
- [15] BEREC, "An Assessment of IP Interconnection in the Context of Net Neutrality," December 2012. [Online]. Available: http://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/1130-an-assessment-of-ip-interconnection-in-t_0.pdf. [Accessed 14 January 2013].
- [16] J. S. Marcus and A. Monti, "Network operators and content providers: Who bears the cost?," *WIK-Consult*, 2011.
- [17] Telegeograph, "International Market Trends," [Online]. Available: http://www.ptc.org/ptc12/images/papers/upload/PTC12_TeleGeograph%20Wkshop.pdf. [Accessed 28 March 2013].

- [18] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide and F. Jahanian, "Internet Inter-domain Traffic," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4, p. 75–86, 2010.
- [19] D. Weller, "Blurring Boundaries: Global and Regional IP Interconnection," GSR Discussion Paper, 2012.
- [20] Netflix, "Announcing the Netflix Open Connect Network," [Online]. Available: <http://blog.netflix.com/2012/06/announcing-netflix-open-connect-network.html>. [Accessed 12 September 2012].
- [21] L. Chapin and C. Owens, "Interconnection and Peering Among Internet Service Providers," 2005 Interisle Consulting Group. [Online]. Available: <http://www.interisle.net/sub/ISP%20Interconnection.pdf>. [Accessed 7 January 2013].
- [22] J. Baldwin, J. Ewert and S. Yamen, "Evolution of the Voice Interconnect," *Ericsson Review* 2, pp. 10-15, 2010.
- [23] Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic," Cisco, May 2012. [Online]. Available: http://www.cisco.com/en/US/netsol/ns827/networking_solutions_white_papers_list.html. [Accessed 30 May 2013].
- [24] GSM Association Press Release, [Online]. Available: <http://www.gsma.com/newsroom/gsma-the-rise-of-connected-devices-will-drive-mobile-operator-data-revenues-past-voice-revenues-globally-by-2018>. [Accessed 28 March 2013].
- [25] M. Scott, D. Elixmann and K. Carter, "The Future of IP interconnection: Technical, Economic, and Public Policy Aspects," *WIK-Consult Report for the European Commission*, January 2008.
- [26] C. Growitsch, M. Scott and W. Christian, "The Effects of Lower Mobile Termination Rates on Retail Price and Demand," *Communications and Strategies*, vol. 80, no. 4th Quarter, pp. 119-140, 2010.
- [27] S. Fabrizi and B. Wertlen, "Roaming in the Mobile Internet," *Telecommunications Policy*, vol. 32, no. 1, pp. 50-61, 2008.
- [28] T. Stühmeier, "Roaming and Investments in the Mobile Internet Market," *Telecommunications Policy*, vol. 36, no. 8, p. 595–607, September 2012.
- [29] BEREC, "International Mobile Roaming Regulation," December 2010. [Online]. Available: [http://www.erg.eu/streaming/BoR%20\(10\)%2058%20International%20Roaming_final_110209.pdf?contentId=547089&field=ATTACHED_FILE](http://www.erg.eu/streaming/BoR%20(10)%2058%20International%20Roaming_final_110209.pdf?contentId=547089&field=ATTACHED_FILE). [Accessed 13 October 2012].
- [30] GSM Association, "IR.34, Inter-Service Provider IP Backbone Guidelines Version 7.0.," 2012.
- [31] AMS-IX, "Annual Report 2012," [Online]. Available: <https://www.ams-ix.net/about/organisation/annual-reports>. [Accessed 21 October 2012].
- [32] GSM Association, "IPX White Paper, Version 1.2," 2007.
- [33] i3 Forum, "Understanding the Different Flavors of IPX," May 2011. [Online]. Available: http://i3forum.org/sites/default/files/i3forum_understanding_the_several_flavors_of_IPX_may_2011_V1.pdf. [Accessed 12 December 2012].
- [34] W. Cutler et al., "Quality of Service (QoS) over the Network-to-Network Interface

- (NNI)," MSF Services & Architectures Working Group, 2010.
- [35] M. Dodd, A. Jung, B. Mitchell, P. Paterson and P. Reynolds, "Bill-and-keep and the Economics of Interconnection in Next-Generation Networks," *Telecommunications Policy*, vol. 33, no. 5-6, pp. 324-337, 2009.
 - [36] M. Carroll and J. C. Tanner, "IPX: New Interconnect Strategies for Operators," Sybase, June 2011. [Online]. Available: http://www.telecomsemea.net/system/files/IPX_Research_Report_Syabse365_0.pdf. [Accessed 10 December 2012].
 - [37] M. Reuver, H. Bouwman, G. Prieto and A. Visser, "Governance of flexible mobile service platforms," *Futures*, vol. 43, no. 9, pp. 979-985, November 2011.
 - [38] D. Clark, J. Wroclawski, K. R. Sollins and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, pp. 462-475, June 2005.
 - [39] Route Views, [Online]. Available: <http://www.routeviews.org/>. [Accessed 21 September 2012].
 - [40] Routing Information Service (RIS), [Online]. Available: www.ris.ripe.net. [Accessed 22 September 2012].
 - [41] CAIDA, [Online]. Available: <http://as-rank.caida.org/>. [Accessed 2 November 2012].
 - [42] DIMES, [Online]. Available: <http://www.netdimes.org/>. [Accessed 28 October 2012].
 - [43] iPlane, [Online]. Available: <http://iplane.cs.washington.edu/>. [Accessed 11 November 2012].
 - [44] A. Antony and H. Uijterwaal, "Routing Information Service RIS Design Note," RIPE NCC, 1999.
 - [45] M. Roughan, W. Willinger, O. Maennel, D. Perouli and R. Bush, "10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems," *IEEE Journal on Selected Areas in Communications*, pp. vol. 29, no. 9, pp. 1810-1821, Oct. 2011.
 - [46] L. Blunk, "Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format," RFC 6396, 2011.
 - [47] B. Donnet and O. Bonaventure, "On BGP Communities," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 55-59, 2008.
 - [48] Y. Rekhter, T. Li and S. Hares, "A Border Gateway Protocol (BGP-4)," RFC 4271, 2006.
 - [49] E. Chen and T. Bates, "An Application of the BGP Community Attribute in Multi-home Routing," RFC 1998, 1996.
 - [50] R. Chandra, P. Traina and T. Li., "BGP Communities Attribute," RFC 1997, 1996.
 - [51] RIPE Database Query, [Online]. Available: <https://apps.db.ripe.net/search/query.html>. [Accessed 5 November 2012].
 - [52] RIPE, "RIR Statistics Exchange Format," [Online]. Available: <ftp://ftp.ripe.net/ripe/stats/RIR-Statistics-Exchange-Format.txt>. [Accessed 9 December 2012].
 - [53] Peering Database, [Online]. Available: www.peerindb.com. [Accessed 1 October 2012].
 - [54] Packet Clearing House, [Online]. Available: www.phc.net. [Accessed 11 October 2012].

- [55] Euro-IX, [Online]. Available: <https://www.euro-ix.net/>. [Accessed 10 October 2012].
- [56] R. Govindan and A. Reddy, "An Analysis of Internet Inter-domain Topology and Route Stability," in *Annual Joint Conference of the IEEE*, p. 850–857, 1997.
- [57] M. Faloutsos, P. Faloutsos and C. Faloutsos, "On Power-law Relationships of the Internet Topolog," *ACM SIGCOMM Computer Communication Review*, vol. 29, no. 4, pp. 251-262, 1999.
- [58] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Transactions on Networking (ToN)*, vol. 9, no. 6, pp. 733-745, 2001.
- [59] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy and G. Riley, "AS Relationships: Inference and Validation," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 37, no. 1, pp. 29-40, 2007.
- [60] L. Subramanian, S. Agarwal, J. Rexford and R. H. Katz, "Characterizing the Internet Hierarchy from Multiple Vantage Points," *IEEE INFOCOM*, 2002.
- [61] G. D. Battistac, M. Patrignani and M. Pizzonia, "Computing the Types of the Relationships between Autonomous Systems," *IEEE INFOCOM*, 2003.
- [62] T. Erlebach, A. Hall and T. Schank, "Classifying Customer-Provider Relationships in the Internet," in *Proceedings of the IASTED*, 2002.
- [63] X. Dimitropoulos, D. Krioukov, B. Huffaker, k. claffy and G. Riley, "Inferring AS Relationships: Dead End or Lively Beginning," in *International Workshop on Efficient and Experimental Algorithms (WEA)*, 2005.
- [64] S. Kosub and B. Hummel, "Acyclic Type-of-Relationship Problems on the Internet: An Experimental Analysis," in *Proceedings of the 7th ACM SIGCOMM Internet Measurement Conference (IMC)*, pp. 221-226, 2007.
- [65] H. Roberts and D. Larochelle, "Mapping Local Internet Control," in *Computer Communications Workshop*, Hyannis, CA, 2011.
- [66] S. Zhou and G. Zhang, "Chinese Internet AS-level topology," *IET Communications*, vol. 1, no. 2, p. 209–214, 2007.
- [67] M. Wählisch and T. Schmidt, "Exposing a Nation-Centric View on the German Internet," in *Passive and Active Measurement Lecture Notes in Computer Science*, Heidelberg, Springer, 2012, p. 200–210.
- [68] Team Cymru Commercial Services, [Online]. Available: <https://www.team-cymru.org/>. [Accessed 28 February 2013].
- [69] B. Huffaker and M. Fomenkov, "Geocompare: A Comparison of Public and Commercial Geolocation Databases," in *Network Mapping and Measurement Conference (NMMC)*, 2011.
- [70] MaxMind's GeoLite Country Database, [Online]. Available: <http://dev.maxmind.com/geoip/legacy/geolite/>. [Accessed 6 January 2013].
- [71] RIPE NCC, "IPv4 Exhaustion," [Online]. Available: <http://www.ripe.net/internet-coordination/ipv4-exhaustion>. [Accessed 15 September 2012].
- [72] M. Xiaoqiao, Z. Xu, B. Zhang, G. Huston, S. Lu and L. Zhang, "IPv4 Address Allocation and the BGP Routing Table Evolution," in *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 1, pp. 71-80, 2005.
- [73] R. Oliveira, B. Zhang and L. Zhang, "Observing the evolution of Internet AS topology,"

- ACM SIGCOMM Computer Communications Review*, vol. 37, no. 4, pp. 313-324, 2007.
- [74] Z. Xu, X. Meng, C. J. Wittbrodt, L. Zhang and S. Lu, "IPv4 Address Allocation and the Evolution of the BGP Routing Table," in *Proceedings of IEEE Computer Communications Workshop (CCW)*, 2003.
 - [75] X. Cai, J. Heidemann, B. Krishnamurthy and W. Willinger, "Towards an AS-to-Organization Map," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet measurement*, pp. 199-205, 2010.
 - [76] J. L. S. Damas, A. Robachevsky and D. Walker, "RIPE Whois Database Query Reference," RIPE NCC, October 2005.
 - [77] RIPE Stats, [Online]. Available: <ftp://ftp.ripe.net/ripe/stats>. [Accessed 5 November 2012].
 - [78] M. Roughan, S. Tuke and O. Maennel, "Bigfoot, sasquatch, the yeti and other missing links: what we don't know about the as graph," in *the 8th ACM SIGCOMM Conference*, p. 325-330, 2008.
 - [79] S. Y. Qiu, P. D. Mcdaniel and F. Monrose, "Toward valley-free interdomain routing," *IEEE ICC*, pp. 2009-2016, 2007.
 - [80] Y. Zhang and M. Tatipamula, "Characterization and design of effective BGP AS-path prepending," in *the 19th IEEE International Conference on Network Protocols (ICNP)*, 2011.
 - [81] MSCI, "The Global Industry Classification Standard," [Online]. Available: <http://www.msci.com/products/indices/sector/gics/>. [Accessed 22 December 2012].
 - [82] Cytoscape, [Online]. Available: <http://www.cytoscape.org/>. [Accessed 19 September 2012].
 - [83] X. Liu and L. Xiao, "A Survey of Multihoming Technology in Stub Networks: Current Research and Open Issues," *Network, IEEE*, vol. 21, no. 3, pp. 32-40, 2007.
 - [84] FICIX, [Online]. Available: www.ficix.fi. [Accessed 3 September 2012].
 - [85] Tampere Region Exchange, [Online]. Available: <http://www.trex.fi/>. [Accessed 7 September 2012].
 - [86] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig and W. Willinger, "Anatomy of a Large European IXP," in *ACM SIGCOMM 2012 Conference*, pp. 163-174, 2012.
 - [87] B. Augustin, B. Krishnamurthy and W. Willinger, "IXPs: Mapped?," in *Proceedings of the 9th ACM SIGCOMM*, pp. 336-349, 2009.
 - [88] CYB Newlog, "ENUM Commercial in Finland +358," 5 October 2006. [Online]. Available: <http://www.itu.int/ITU-D/cyb/newslog/ENUM+Commercial+In+Finland+358.aspx>. [Accessed 22 December 2012].
 - [89] ENUM, "RIPE ENUM Working Group," [Online]. Available: www.enumdata.org. [Accessed 28 February 2013].
 - [90] N. Gelvanovska and J. Alden, "Coexistence of Traditional and IP Interconnection," GSR Discussion Paper, 2009.
 - [91] GSM Association White Paper, "Packet Voice Interworking for Mobile Service Providers," July 2008. [Online]. Available: <http://www.gsma.com/technicalprojects/wp-content/uploads/2012/05/packetvoiceinterworkwp-1.pdf>. [Accessed 7 November 2012].

- [92] M. Kende, "Overview of recent changes in the IP interconnection ecosystem," Analysys Mason, 23 January 2011.
- [93] A. B. Bondi, "Characteristics of Scalability and Their Impact on Performance," in *Proceedings of the 2nd International Workshop on Software and Performance*, Ottawa, Ontario, Canada, 2000.
- [94] RIPE NCC, [Online]. Available: <http://www.ripe.net/>. [Accessed September 27 2012].



Appendix A – The flow chart for expanding the filter list of ASes and prefixes in longitudinal analysis




FIGURE 25 - FLOW CHART FOR LONGITUDINAL ANALYSIS (FIRST PART)




FIGURE 26 - FLOW CHART FOR LONGITUDINAL ANALYSIS (SECOND PART)

Appendix B – The hierarchical layout transit links




FIGURE 27 - HIERARCHICAL LAYOUT OF TRANSIT RELATIONSHIPS